

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
5 July 2001 (05.07.2001)

PCT

(10) International Publication Number
WO 01/48637 A1

(51) International Patent Classification⁷: **G06F 17/30**

(21) International Application Number: **PCT/US00/34413**

(22) International Filing Date:
18 December 2000 (18.12.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
09/474,929 29 December 1999 (29.12.1999) US

(71) Applicants: INCYTE GENOMICS, INC. [US/US]; 3160 Porter Drive, Palo Alto, CA 94304 (US). STANFORD UNIVERSITY [US/US]; Office of Technology Licensing, Suite 350, 900 Welch Road, Palo Alto, CA 94304-1850 (US).

(72) Inventors: WALKER, Michael, G.; 1050 Borregas Ave., #80, Sunnyvale, CA 94089 (US). WANG, James, Z.; 672 Chiquita #3, Mountain View, CA 94041 (US). GILADI, Eldar, Y.; 1063 Colorado Ave., #2-H, Palo Alto, CA 94303 (US).

(74) Agents: TRINIDAD, Arriola-Kern et al.; Fenwick & West LLP, Two Palo Alto Square, Palo Alto, CA 94306 (US).

(81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.

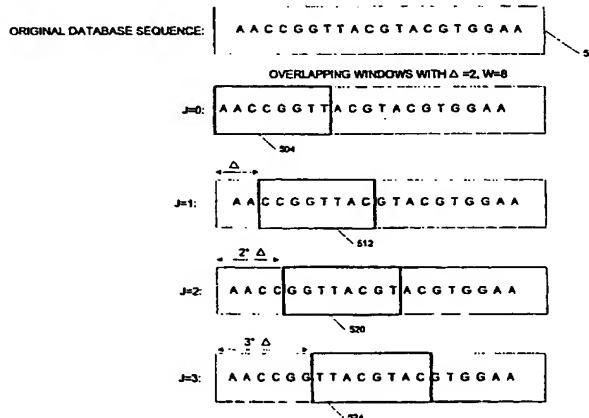
(84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

Published:

- With international search report.
- Before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments.

[Continued on next page]

(54) Title: SEQUENCE DATABASE SEARCH WITH SEQUENCE SEARCH TREES



WO 01/48637 A1

(57) Abstract: A method and system for generating and searching a tree-structured index of window vectors that represent database sequences comprise a window vector generation module, a tree-structured index generation module, a query sequence partitioning module, and a retrieval component. The window vector generation module partitions a database sequence into a plurality of overlapping windows. Each window has a fixed length W comprising a fixed number of nucleotides, and the offset among windows is determined by a parameter Δ . The window vector generation module then maps each database sequence window into a window vector. The database sequence window vector indicates the frequency of appearance of each k-tuple in the corresponding database sequence window. The tree-structured index generation module then generates a tree-structured index using the database sequence window vectors. The query sequence partitioning module partitions a query sequence into a plurality of windows and maps each query sequence window into a query sequence window vector. Each query sequence window vector is then compared against the tree-structured index to locate the database sequences that are similar to the query sequence. The list database sequences that are similar to the sequence is then returned as the result of the search.



For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SEQUENCE DATABASE SEARCH WITH SEQUENCE SEARCH TREES

5

GOVERNMENT SUPPORT

Work for the following disclosed invention was supported in part by the National
10 Science Foundation under Grant No. _____. The Government may have certain rights in
the invention.

BACKGROUND

A. Field of the Invention

The present invention relates generally to storage, indexing, and retrieval of
15 database sequences, and more particularly to a method and system for generating and
searching a tree-structured index of window vectors that represent database sequences.

B. Background of the Invention

Sequence-similarity finding programs identify sequences in DNA and
protein databases that are similar to a query sequence. Because of the recent explosion in
20 the amount of DNA sequence information available in public and private databases as a
result of the human genome project and other large-scale DNA sequencing efforts, such
sequence-similarity finding programs have become increasingly important in modern
biology.

Generally, there are two classes of sequence-similarity searching programs:
25 global comparison methods (e.g. Needleman and Wunsch), and local comparison methods,
e.g., the FASTA method (Pearson and Lipman), and the BLAST method (Altschul). Global
comparison methods have a high degree of accuracy but are extremely slow. Local
comparison methods such as FASTA and BLAST identify candidate similar sequences
based on shared k-tuples and therefore are faster than global methods. However, local
30 comparison methods are less accurate, i.e. they provide results which are not similar.
Moreover, with both the global and local comparison methods, the computational

complexity of such methods usually increases linearly with the size of the number of sequences to be searched. This is due in part to the fact that most prior art methods search all, or at least a very large part, of the sequence database. To improve the searching time, several prior art methods have created sequence databases using clusters of sequences, 5 including tree-structured indexes. However, many of these methods use conventional sequence alignment methods such as BLAST and FASTA to determine pairwise distances between the sequences. Thus, such cluster and tree-structured methods are also limited by the speed of the alignment methods described above.

10 Accordingly, what is needed is a system and method for organizing and searching database sequences that is fast and efficient, and at the same time provides a high degree of accuracy, that is, one that identifies sequences similar to a query sequence.

SUMMARY OF THE INVENTION

15 The present invention overcomes the limitations of conventional sequence-similarity-searching programs by using window vectors that represent database sequences in a sequence storage and retrieval system. The window vectors associated with the database sequences are organized into a tree-structured index for faster and more efficient searching of the database sequences. A query sequence is used to search the tree-structured index for database sequences that are similar to the query sequence.

20 In one embodiment, each database sequence is partitioned into a plurality of overlapping windows or fragments of fixed length. Each database sequence window has a fixed length W and the degree of offset between windows is determined by a parameter Δ . Each database sequence window comprises a subsequence of elements from the database sequence beginning at position $j*\Delta$, from the start of the database sequence, and ending at position $j*\Delta+W$, where $j = 0, 1, 2, \dots, (L-W)/\Delta$. In other words, each database sequence 25 window has a fixed length W which is advanced down the length of the original database sequence every Δ elements.

30 For each database sequence window, a database sequence window vector is computed. A database sequence window vector represents the occurrence of each k-tuple in the database sequence window. In one embodiment, the database sequence refers to a DNA or protein sequence, and the occurrence of each k-tuple in the DNA or protein sequence window is represented by a vector of length 4^k . Each position in the vector represents a unique k-length sequence (i.e. k-tuple). If a k-tuple occurs more than one time in a database

sequence window, either the number of times that the k-tuple occurs may be recorded in the corresponding position in the vector or the value "1" may be recorded in the corresponding position in the vector to indicate that the k-tuple occurs at least once in the database sequence window. If a k-tuple does not occur in a database sequence, a zero in the 5 corresponding position in the vector may be used to indicate that the k-tuple does not occur in the sequence. In one embodiment, database sequence window vectors are stored in a tree-structured index to reduce searching time.

Database sequences are retrieved using a query sequence. In one embodiment, a query sequence is partitioned into a plurality of windows. For each query sequence 10 window, a query sequence window vector is computed. A query sequence window vector represents the occurrence of each k-tuple in the query sequence window. Each query sequence window vector is compared against the tree-structured index of the database sequence window vectors to locate the nearest neighbors of the query sequence, i.e. database sequences similar to the query sequence. In other words, for each query sequence 15 window vector, the tree-structured index is traversed from a root node of the tree to a terminal node which contains the nearest neighbor for that query sequence window vector. In one embodiment, the list of sequences with at least one significant window hit is returned.

In one embodiment, the present invention is used to generate and search a tree- 20 structured index of window vectors that represent biological database sequences. Each database sequence may represent a DNA sequence comprising a fixed number of nucleotides. The DNA database sequence is then partitioned into a plurality of overlapping windows. Each DNA database sequence window has a fixed length W comprising a fixed number of nucleotides, and the degree of offset among windows is determined by a 25 parameter Δ . Each DNA database sequence window is then mapped into a database sequence window vector. The DNA database sequence window vector indicates the frequency of appearance of each k-tuple in the corresponding DNA database sequence window. A tree-structured index is then generated using the DNA database sequence window vectors. To search the tree-structured index, a query sequence (e.g. a DNA 30 sequence) is partitioned into a plurality of windows. Each query sequence window is then mapped into a query sequence window vector. Each query sequence window vector is then compared against the tree-structured index to locate the DNA database sequence window

vectors which are closest to the DNA query sequence window vector. The list of DNA database sequences that is similar to the DNA query sequence is then returned.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a system in accordance with the present invention.

5 Fig. 2 is a functional block diagram of the present invention.

Fig. 3 is a flow chart of the steps of a preferred embodiment of the present invention.

10 Fig. 4 is a flow chart of the steps of a preferred embodiment for generating window vectors in accordance with the present invention.

Fig. 5 is an illustration of a preferred embodiment for partitioning a database sequence in accordance with the present invention.

15 Fig. 6 is a graph showing the results of computations performed to monitor the sensitivity of the present invention as a function of the offset size and the tuple size.

Fig. 7 is a flow chart of the steps of a preferred embodiment for generating a tree-structured index of the database sequences in accordance with the present invention.

20 Fig. 8 is a flow chart of the steps of a preferred embodiment for partitioning and using a query sequence to search a tree-structured index in accordance with the present invention.

Fig. 9 is a graph showing the true positive rate as a function of the window size.

25 Fig. 10 is a graph showing the true positive rate as a function of window size.

Fig. 11 is a graph showing the false positive rate as a function of the window size and threshold.

Fig. 12 is a graph showing the false positive rate as a function of the window size and threshold.

25 Fig. 13 is a graph showing the ratio of the number of false positive pairs to the number of true positive pairs.

Fig. 14 is a graph showing the time to generate a tree-structured index and the time to search it as a function of window size.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

A. System Overview

30 Referring now to Figure 1, there is shown a block diagram of software and hardware components for implementing one embodiment of the present invention 100. Processor 102

is a conventional engineering workstation or other computer processor and may also be coupled to other processors accessible over conventional communications channels or buses (not shown). Processor 102 is conventionally coupled to storage 104 which may be a magnetic disk storage, a CD storage unit, or other conventional computer data storage unit.

5 Storage 104 may also be coupled to other storage units accessible over conventional communications channels or buses (not shown). Storage 104 includes sequence database 106. Sequence database 106 contains a plurality of sequences, also referred to herein as "database sequences", to be searched. In a preferred embodiment, the database sequences represent biological sequences, such as DNA or protein sequences, that are represented in electronic or computer readable form. The present invention however is not limited to DNA sequences, but may be used with other types of data.

10

Processor 102 is also conventionally coupled to memory 108 which is a random access memory (RAM) unit or other conventional computer memory. Items in memory 108 may alternatively be stored in storage 104 and accessed by processor 102 when required.

15 Memory 108 may comprise various software products including window vector generation 110, tree-structure index generation 112, tree structure index 114, query sequence partitioning 116, and retrieval 118 program components or modules as discussed below. Input 101 comprises conventional input devices such as a keyboard, mouse, track ball, or touchscreen. A conventional display unit 120 may also be conventionally coupled to

20 processor 102.

In a preferred embodiment, the present invention may be implemented on any platform, operating system, and user interface of sufficient power and flexibility, such as: XWindows/MOTIF; Sun/OS SunView; Microsoft Windows, VAX/VMS, and the like, and may be operated in a distributed data / distributed processor environment, if desired. In a preferred embodiment, the invention is capable of handling a very large data corpus, e.g., containing over 10 million sequences.

25

Referring now to Figures 2 and 3, there are shown a functional model of the present invention for generating window vectors and identifying database sequences similar to an input query sequence and a flow chart of the steps of a preferred use of the present invention. In this embodiment, the window vector generation module 110 generates 302 a plurality of window vectors associated with each database sequence. Preferably, the window vector generation module 110 partitions each database sequence into a plurality of windows and then encodes each window into a window vector. Preferred embodiments of

the functional operation of the window vector generation module 110 are described in more detail below with reference to Figure 4. The tree-structure index generation module 112 then generates 304 a tree-structured index 114 of the window vectors, preferably using Tree-Structured Vector Quantization (TSVQ). A preferred embodiment of the functional 5 operation of the tree-structure index generation module 112 is described below in more detail with reference to Figure 7. A query sequence 206 is input 306 into the system, and the query sequence partitioning module 116 partitions 308 the query sequence into a plurality of windows. In a preferred embodiment, the query sequence partitioning module 116 then maps or encodes 310 each query sequence window into a query sequence window 10 vector. A preferred embodiment of the functional operation of the query sequence partitioning module 116 is described below in more detail with reference to Figure 8. Finally, the retrieval module 118 then uses the query sequence window vector to search 312 the tree-structured index 114 to identify nearest neighbor window vectors for the query sequence, i.e. database sequences which are similar to the query sequence. A preferred 15 embodiment of the functional operation of the retrieval module 118 is described below in more detail with reference to Figure 8.

Thus, the present invention provides vector quantization and clustering methods for organizing databases for rapid and efficient searching. The present invention is based upon a database sequence search tree where the database sequences are represented using 20 database sequence window vectors organized into a tree-structured index. Window vectors represent the occurrences of each k-tuple of interest in a database sequence, or a query sequence, as described in more detail below. A k-tuple, as used herein, is a fixed length ordered sequence of elements and is well known to one skilled in the art of sequence-similarity searching. The use of window vectors advantageously increases the sensitivity of 25 the present invention as described in more detail below. The database sequence window vectors are then used to generate a tree-structured index representing the database sequences. The tree-structured index decreases the searching time of the present invention by allowing the present invention to search only a portion of the database sequences to find 30 similar sequences. Thus, whereas the searching time required for prior art sequence-similarity searching methods increases linearly with the size of the database, the searching time required for the present invention increases logarithmically with the size of the database.

B. Window Vector Generation

Referring now to Figure 4, there is shown a high-level flowchart of a preferred embodiment of the steps for generating database sequence window vectors in accordance with the present invention. Each database sequence comprises a plurality of elements, L, and is partitioned 402 into a plurality of windows or fragments. Preferably, a "sliding window scheme" is used to partition each database sequence into overlapping windows or fragments. In other words, each database sequence window has a fixed length W and the degree of offset between windows is determined by a parameter Δ . In one embodiment, the length of the window W refers to the number of DNA nucleotides or protein amino acids in the database sequence. Each database sequence window consists of subsequences of the original database sequence beginning at position $j*\Delta$, in the original database sequence, and ending at position $j*\Delta + W$, where $j = 0, 1 \dots (L-W)/\Delta$.

Referring now to Figure 5, there is shown an example of partitioning a database sequence in accordance with the sliding window scheme. For example, assume an original database sequence 502 comprises a DNA sequence of 20 nucleotides ($L=20$) where each nucleotide is represented by one of the letters A,C,G, or T. The original database sequence 502 is partitioned into a plurality of windows using the parameters $W=8$ and $\Delta = 2$. Accordingly, for $j=0$, window 504 is created. For $j=1$, window 512 is created. For $j=2$, window 520 is created. For $j=3$, window 524 is created. The process is continued until all the database sequence windows for the database sequence have been created, that is, until $j=(L-W)/\Delta$ (until $j=6$ in this example). Thus, for the example in Figure 5, seven overlapping windows will be generated for the original database sequence 502.

As can be seen in Figure 5, the offset between each window is a multiple of Δ . For example, the beginning of window 512 is offset from window 504 by Δ resulting in an overlap of the nucleotides CCGGTT. The beginning of window 520 is offset from window 504 by $2*\Delta$ and from window 512 by Δ . Therefore, between windows 520 and 504 there is an overlap of nucleotides GGTT, and between windows 520 and 512 there is an overlap of nucleotides GGTTAC. Thus, each window has a fixed length W which is advanced down the length, L, of the original database sequence every Δ elements. In other words, each window runs from $j*\Delta$ to $j*\Delta + W$, where $j = 0, 1 \dots (L-W)/\Delta$. The sliding window scheme is used to generate a plurality of overlapping windows until the end of the

DNA database sequence is reached. In a preferred embodiment, values for Δ , W, and L are in the range $0 \geq W \geq 25$; $5 \geq \Delta \geq 50$; and $20 \geq L \geq 1000$.

Partitioning the database sequences into overlapping windows advantageously increases the sensitivity of the searching algorithm. Referring now to 5 Figure 6, there is shown the results of computations performed to monitor the sensitivity of the present invention as a function of the offset size, Δ , and the tuple size, k. The ordinate of Figure 6 is the percentage of overlaps detected and the abscissa is the tuple size used in the encoding. Each graphed line corresponds to a different offset size Δ in the sliding window encoding scheme. For example, the solid line corresponds to encoding with no 10 windows, while the dashed line, $\Delta=50$, corresponds to the use of a single offset. Each graphed line in the figure shows the percentage of overlaps correctly detected by the algorithm as a function of the tuple size k. Only overlaps of size greater than or equal than 50 were considered. Thus, as illustrated in Figure 6, the use of the "sliding window" scheme increases the sensitivity of the present invention. Moreover, as the tuple size k 15 increases, the sensitivity of the present invention also increases.

Referring again to Figure 4, each database sequence is partitioned into a plurality of overlapping windows and each database sequence window is mapped into a window vector. One skilled in the art will realize that the order in which the database sequences are partitioned into a plurality of windows and the order in which the database 20 sequence windows are mapped into database sequence window vectors is not essential to the invention. In other words, a single database sequence may be partitioned into a plurality of windows and afterwards, each database sequence window may then be used to compute the database sequence window vector corresponding to each database sequence window. In another embodiment, the present invention may first partition all the database sequences 25 into a plurality of database sequence windows, and afterwards compute the corresponding database sequence window vectors. In yet another embodiment, the present invention may partition a database sequence into a first database sequence window and then compute the database sequence window vector for the first database sequence window. The present invention may then partition the database sequence into a second database sequence 30 window and compute the database sequence window vector for the second database sequence window. Thus, one skilled in the art will realize that many options are available for the sequence of steps for partitioning each database sequence into a plurality of

overlapping windows and computing the database sequence window vectors corresponding to each database sequence window.

Referring back to Figure 4, there is shown one embodiment of a method for generating database sequence window vectors for a database sequence. After each database sequence has been partitioned 402 into a plurality of overlapping windows, each database sequence window is then mapped into a window vector. The database sequence window vectors are generated by first determining for each window 404 the frequency of appearance of each k-tuple in a specified database sequence window. A database sequence window vector is then generated 406 for the specified database sequence window. The database sequence window vector indicates the frequency of each k-tuple in that corresponding database sequence window. For a sequence, a k-tuple is a fixed length ordered sequence of elements.

In general, the number of possible k-tuples of length k is N^k , where N is the number of unique identifiers from which a sequence can be created and k is the tuple length. For example, from an alphabet of four letters, as in DNA for example, the total number of possible k-tuples is 4^k . For example, in a sequence generated from the DNA alphabet consisting of {A, C, G, T}, there are 16 possible ordered k-tuples of length k=2: {AA, AC, AG, AT, CA, CC, CG, CT, GA, GC, GG, GT, TA, TC, TG, TT}. In the case of protein sequences, there are 20 letters from which a protein sequence can be created, giving 20^k possible k-tuples. Alternate amino-acid alphabets with fewer letters would yield smaller numbers of k-tuples for a fixed k.

Referring to a DNA sequence as an example, the occurrence of each k-tuple in a DNA sequence may be represented by a window vector of length 4^k . For example, the DNA sequence AACTT contains 2-tuples AA, AC, CT, and TT. These occurrences of 2-tuples may be represented by the 16-position window vector {1100000100000001}. This window vector corresponds to the ordered representation {AA, AC, AG, AT, CA, CC, CG, CT, GA, GC, GG, GT, TA, TC, TG, TT}. Thus, the 1 in the first position corresponds to the presence of the k-tuple AA, the 1 in the second position corresponds to the presence of the k-tuple AC, the 1 in the eighth position corresponds to the presence of the k-tuple CT, and the 1 in the sixteenth position corresponds to the presence of the k-tuple TT. If a particular k-tuple is not observed in a sequence, then a zero may be placed in the corresponding position in the window vector, as is the case for AG, AT, etc.... If a k-tuple occurs more than one time in a sequence, either the number of times that the k-tuple occurs

may be recorded in the corresponding position in the window vector or the value “1” may be recorded in the corresponding position in the window vector to indicate that the k-tuple occurs at least once in the sequence. The k-tuple size may be of any length, but for DNA and protein database sequences, k is preferably in the range of $2 \leq k \leq 10$. Thus, window
5 vectors provide fast and accurate identification of similar sequences with acceptable space requirements.

In one embodiment, a sparse window vector representation may be used to generate the window vectors. Sparse window vector representation allows a user to store only the non-zero entries by storing (index, value) pairs rather than storing the whole vector.
10 For example, consider the DNA sequence ACCGTTGACACCC and suppose that $k=2$, $\Delta=2$, and $W=5$. On a disk or larger memory space, all tuples and an index to the beginning of each window are stored. For example, for the above DNA sequence, the following tuple information is stored: AC,CC, CG,GT,TT,TG,GA,AC,CA,AC,CC,CC. In a smaller memory space, such as a RAM, the windows are stored as sparse window vectors such as
15 $\{(AC, 1), (CC,1), (CG,1), (GT,1)\}$. Accordingly, when sparse window vectors are used, sparse window vector computations should also be used. In other words, only non-zero entries are used for computations. If an entry is zero in both vectors, it is zero when both vectors are subtracted and hence, no arithmetic needs to be done on entries that are zero. Thus, sparse window vector representation advantageously decreases the amount of
20 memory required for storing the vector windows. Additionally, by using sparse window vector computations, the present invention provides for faster and more efficient computation.

C. Tree-Structure Index Generator

25 The tree-structured index or database sequence search tree (“SST”) organizes the window vectors that represent the database sequences to enable fast and efficient searching of the database sequences. In a preferred embodiment, a binary-tree structured index is used. Alternatively, the tree-structured index may use 3-way branching, 4-way branching, or an arbitrary number of branches at each branch node in the tree. The binary-
30 tree structured index provides for fast searching but requires the largest amount of space. Thus, the branching factor may be selected to best suit the constraints of the available computer hardware.

The tree-structured index is generated by recursively searching the database sequence window vectors for clusters that provide binary (or higher order) partitions. A variety of methods are available for finding such clusters and building the tree-structured index as would be known to one skilled in the art. In one embodiment, a Tree-Structured Vector Quantization (TSVQ) method using k-means clustering is used to generate the tree-structured index. Referring now to Figure 7, there is shown a high-level flowchart of a preferred embodiment of the steps for generating a tree-structured index of the database sequence window vectors in accordance with the present invention. First, two centroids and their corresponding partitions of the data are determined 702. In a preferred embodiment, this is done by choosing and designating two points, a_0 and b_0 , in the k-dimensional space as centroids. As used herein, a centroid represents a group of database sequence window vectors. The set of all database sequences closest to each centroid comprises a partition of the k-dimensional space. For example, in a 2-dimensional subspace, the database sequences are partitioned into two disjoint subsets, A and B, by assigning to set A, the sequences closer to a_0 , and assigning to set B all the other sequences. In other words, for each database sequence, the distance of each sequence from each centroid is calculated and the sequence is assigned to the partition corresponding to the nearest centroid, i.e., the centroid, a_0 and b_0 , which is a minimum distance from the sequence. A new centroid is then determined for each partition by calculating the sum of the distances of all the sequences in the partition from the centroid using the following equations:

$$a_1 = \sum_{v \in A} v / |A| \quad b_1 = \sum_{v \in B} v / |B|$$

where A is the set of vectors in Group A, B is the set of vectors in group B, v is a vector, and $\sum_{v \in A}$ is the sum of all vectors belonging to set A.

The steps for determining the two centroids are repeated until convergence, that is, until the change in the sum of the distances at each iteration is less than a small threshold, ϵ , or no vectors change partition. In one embodiment, the following equation is used to determine convergence:

$$||a_0 - a_1|| < \epsilon$$

where ϵ is preferably a small positive number such as $\epsilon = 10^{-10}$.

If there is no convergence, then a_0 is set to a_1 and b_0 is set to b_1 and step 702 is repeated until there is convergence. For each partition generated, two centroids and their corresponding partitions are determined 704 using the steps described above. These steps 5 are repeated 706 until a stopping criterion is met, that is until the number of training sequences in the partitions is less than a user-specified threshold, or the maximum number of terminal nodes is reached. Each of the terminal nodes is a partition of the k-dimensional space and contains a set of sequences that are nearest neighbors to the centroid of that node. In a preferred embodiment, 16-bit finite precision arithmetic is used to compute the 10 centroids. However, one skilled in the art will realize that floating point arithmetic may also be used to compute the centroids.

One skilled in the art will realize that the centroids may be computed using a sample of the windows in the database rather than computing centroids for all the window vectors. For example, in one embodiment, a sample of 500 sequences to estimate two 15 centroids may provide an improvement in the speed for generating a tree-structured index while still providing a relatively low cost in the error rate. Alternatively, one skilled in the art will realize that where higher accuracy is required, a larger sample of sequences for estimating two centroids is preferable. Thus, the use of sampling in generating a tree-structured index enables one to reduce the false negative rate by generating two trees, using 20 different samples, and taking the union of the database sequences in the terminal nodes in which the query sequence falls in the two trees. Both the generation and search of both trees may be done in parallel thereby increasing efficiency and speed. Additionally, in one embodiment, generation of the tree-structured index is completed the first time a leaf size is smaller than 100.

25

D. Query Sequence and Database Sequence Retrieval

A query sequence is used to search the tree-structured index to find the database sequences that represent the nearest neighbors of the query sequence. In one embodiment a query sequence refers to a biological sequence, such as a DNA or protein 30 sequence, for which one wishes to find similar sequences in the sequence database. Referring now to Figure 8, there is shown a high-level flowchart of a preferred embodiment of the steps for partitioning and using the query sequence to search the tree-structured index. Each query sequence is first partitioned 802 into a plurality of windows. Each window of

the query sequence has a fixed length W which corresponds to the length of the database sequence windows. In one embodiment, the query sequence may be partitioned into non-overlapping windows. In another embodiment, the query sequence may be partitioned into overlapping windows where the degree of offset among windows is determined by the parameter Δ as discussed above with respect to window vector generation. Preferably, the query sequence is partitioned using either value for Δ : $\Delta = W$ or $\Delta = W/2$. Each query sequence window is then mapped into a window vector. A query sequence window vector represents the occurrence of each k-tuple in the query sequence window.

Each query sequence window vector is then used to search the tree-structured index by choosing the node whose centroid is the minimum distance between the query sequence window vector and the centroid of that node. The list of database sequences with at least one significant window hit is returned. In other words, the database sequence window vectors in a leaf which are closer to query sequence window vectors than to a given threshold are returned. Additional processing such as alignment of the query sequence to the database sequence with one of the standard alignment tools is also possible. In other words, one can use the Smith-Waterman algorithm to align windows to obtain a better score.

Referring again to Figure 8, the steps for searching a binary tree-structured index are shown. A search of the tree-structured index begins at the root node of the tree. At each non-terminal node in the tree, each branch is represented in the tree by its centroid. The distance between each centroid and the query window is determined 804. In a preferred embodiment, the distance, d , between two windows, such as a database sequence window, W_1 , and a query window, W_2 , is determined by the following equation:

$$d(W_1, W_2) = || X(W_1) - X(W_2) ||$$

where $X(W)$ refers to the window vector of the window.

The branch whose centroid is the minimum distance from the query sequence window is selected 806. Steps 804 and 806 are repeated through the tree, progressing from the root note towards the terminal nodes, until a terminal node is reached 808. The database sequences in the selected terminal node are the nearest neighbors of the query sequence window. The database sequences representing the nearest neighbors of the query sequence window are then retrieved 810 from the sequence database. In other words, for each query

sequence, all database sequences which have a window in common with the query sequence are outputted as a resulting hit.

In another embodiment, the query sequence may be used to search only a portion of the tree-structured index, i.e. a subtree. For example, when searching the tree-
5 structured index using the query sequence window, all query sequence windows may be used to search against a top portion of the tree-structured index. For example, all of the windows of a query sequence may be used to search only the first nine levels of the tree-structured index. Such an embodiment would generate $2^9 = 512$ groups of windows, where each group contains query sequence windows for the search completed in a different leaf of
10 the top subtree. In such an embodiment, each group is searched against the subtree where it hits. With this embodiment, disk access occurs only when the data associated with each subtree is loaded into memory. One skilled in the art will realize that this embodiment may use an arbitrary number of levels for the subtree and that parallel searching of the tree-structured index may be utilized. For example, after the top subtree is processed, lower
15 level subtrees may be processed using different processors. Thus, the present invention provides a faster and efficient method for searching the tree-structured index.

As the number of database sequences increases, it may become difficult to store all the window vectors in memory. In one embodiment, all vectors are stored on a disk and loaded into RAM as needed. In another embodiment, a caching strategy is used.
20 Due to the in-order nature of the tree-structure generation, a FIFO cache may be used. In such embodiment, a refresh mechanism in which a hit in the cache also moves the vector to the front of the queue may also be used. Additionally, the window vector file should be reordered such that all vectors associated with a certain subtree are contiguous.

25 E. Performance Results

Tables 1, 2, and 3 below show the performance results of the present invention in comparison with the prior art method BLAST. In the computations, a 1.552KB piece of genomic DNA was fragmented several times using a Poisson process with average $\Lambda = 300\text{bp}$. From the pool of fragments, three sets of 30675, 61350, and 122700 sequences,
30 respectively, were generated. The present invention and the BLAST2 prior art method were used to determine which of the database sequences overlapped. For the computation using the present invention, a window size of length 50 was used and the offset step $\Delta=5$. The

computation was repeated for several values of the tuple size k, i.e. k = 3, 4, 5, 6. Tables 1, 2, and 3 show the results of the present invention on each of these sets.

Table 1

K	Total	Tree Generation	Query	TPR	FNR	TNR	FPR
3	00:10:12	00:06:56	00:03:16	0.893	0.107	9.9998e-01	2.092e-05
4	00:11:37	00:07:24	00:04:13	0.932	0.068	9.9999e-01	1.069e-05
5	00:13:52	00:08:24	00:05:28	0.956	0.044	9.9999e-01	6.642e-06
6	00:23:04	00:13:04	00:10:00	0.963	0.037	9.9999e-01	5.284e-06

5 Table 1 shows the results of the present invention with 30,675 sequences.

The time for generating the tree-structured index, for tree search, and the total time are displayed for k=3 to 6. The true positive rate (TPR), false negative rate (FNR), true negative rate (TNR), and false positive rate (FPR) are also displayed.

10

Table 2

K	Total	Tree Generation	Query	TPR	FNR	TNR	FPR
3	00:22:33	00:15:36	00:06:57	0.886	0.114	9.9999e-01	1.842e-05
4	00:26:11	00:17:37	00:08:34	0.926	0.074	9.9999e-01	9.805e-06
5	00:36:49	00:19:00	00:17:49	0.953	0.047	9.9999e-01	6.222e-06
6	00:45:53	00:26:37	00:19:16	0.960	0.040	9.9999e-01	5.018e-06

15 Table 2 shows the results of the present invention with 61,350 sequences.

The time for generating the tree-structured index, for tree search, and the total time are displayed for k=3 to 6. The true positive rate (TPR), false negative rate (FNR), true negative rate (TNR), and false positive rate (FPR) are also displayed.

Table 3

K	Total	Tree Generation	Query	TPR	FNR	TNR	FPR
3	00:46:56	00:31:52	00:15:04	0.881	0.119	9.9999e-01	1.529e-05
4	00:50:50	00:32:50	00:18:00	0.924	0.076	9.9999e-01	8.404e-06

5	01:08:22	00:35:21	00:33:01	0.951	0.049	9.9999e-01	5.950e-06
6	01:34:50	00:52:21	00:42:29	0.957	0.043	1.0000e-00	4.793e-06

Table 3 shows the results of the present invention with 122,700 sequences. The time for generating the tree-structured index, for tree search, and the total time are displayed for k=3 to 6. The true positive rate (TPR), false negative rate (FNR), true negative rate (TNR), and false positive rate (FPR) are also displayed.

Searching the database sequences using the BLAST2 prior art method took 02: 03: 09, 07: 27: 47, 28: 10: 25 for 30,675, 61,350, and 122,700, respectively. Thus, the BLAST2 method yields an average time to bin each sequence of 0.2409, 0.4379, 0.8266 seconds. Thus, with the BLAST2 method there is a linear increase in time per sequence. Thus, as can be seen from the data, the search time required for the present invention is virtually independent of the number of database sequences and is for 122700 sequences 0.00074, 0.0088, 0.0161, and 0.0208 for k = 3, 4, 5, 6, respectively. Therefore, for 122,700 sequences the speed-up compared to BLAST2 varies between: 93 for k=3 to 39 for k=6. The search and construction time per sequence using the present invention is 0.0230, 0.0246, 0.0334, and 0.0464 for k = 3, 4, 5, 6, respectively. Therefore, for 122700, the speed-up compared to BLAST2 is 35 for k=3 to 17 for k=6.

F. Determination of Optimal Parameters

Experiments were performed to determine the optimal values for the threshold T, and the window size, W. The results described below were obtained with data with a 5% error rate. Referring now to Figure 9, there is shown a graph indicating the true positive rate (TPR) as a function of the window size, W, and the distance threshold T. Cross-sections of this surface, for fixed values of the window size W are presented in Figure 10. Both figures indicate that the TPR improves as the window size decreases and as the tolerance T increases. Moreover, Figure 9 shows that the TPR has an asymptotic value as T increases. The asymptotic value increases with a decreasing window size.

The FPR also increases with the tolerance T as indicated by Figures 11 and 12 in which the $\log_{10}(FPR)$ is presented. On a logarithmic scale, the growth of the FPR is essentially linear, and the growth rate (the slope) is independent of the window size. Thus, the FPR grows exponentially as a function of T.

In a preferred embodiment, the window size W and the distance threshold T are chosen such that the TPR is large while the number of false positive pairs remain relatively small. Figure 13 indicates the ratio of the number of false positive pairs to the number of true positive pairs returned by the present invention. Thus, as can be seen from 5 Figure 13, for each window size, this ratio increases exponentially with the threshold T. In a preferred embodiment, the parameters W and T are chosen such that their ratio is bounded by 5. Among all possible values for these parameters, the values for which the true positive rate is maximized is desired. The values W=30 and T=15 has a ratio of about 2.5 and a TPR around .95, and the values W=35 and T=20 has a ratio less than 5 and a TPR around 10 .925.

Referring now to Figure 14, there is shown a graph indicating the time to generate a tree-structured index and the time to search it. The time for generating the tree-structured index increases with the window size. The trend for the search time is reversed. To explain this, one can see that the number of windows in the sequence database N is 15 about $N=n/\Delta$, where n is the total number of nucleotides in the database and Δ is the offset parameter. Therefore, N is independent of the window size. However, as the length of a window, W, increases, $X(W)$ has more non-zero entries because W contains more tuples. It follows that the time to compute the distance of $X(W)$ has more non-zero entries because W 20 contains more tuples. Accordingly, the time to compute the distance of $X(W)$ to a centroid increases with the length W, and which in turn increases the overall computation time in generating the tree-structured index. In a preferred embodiment, the number of query windows M per query is $M=m/M$ where m is the number of nucleotides in a query. Thus, as the window size increases, the number of query windows decreases so that the time per query also decreases.

We claim:

1 1. In a computer having a processor and storage, the storage comprising a
2 plurality of database sequences, a computer-implemented process for selecting database
3 sequences, the process comprising:

4 generating a plurality of window vectors representing the database sequences;

5 generating an index of the window vectors; and

6 in response to receiving an input query sequence, searching the index using the
7 input query sequence to select database sequences similar to the query
8 sequence.

9 2. The computer implemented process of claim 1 wherein generating a plurality
10 of window vectors further comprises:

11 partitioning each database sequence into a plurality of overlapping windows ;

12 and

13 generating a window vector for each window.

14 3. The computer implemented process of claim 2 wherein each database
15 sequence comprises a finite number of elements, L, and each window begins at position $j * \Delta$ and ends at position $j\Delta + W$, where $j = 0, 1 \dots (L-W)/\Delta$.

17 4. The computer implemented process of claim 1 wherein generating the index
18 further comprises:

19 searching the window vectors for clusters providing partitions.

20 5. The computer implemented process of claim 1 wherein the tree-index is a
21 tree-structured index.

22 6. The computer implemented process of claim 1 further comprising partitioning
23 the input query sequence into a plurality of non-overlapping windows.

24 7. The computer implemented process of claim 1 further comprising partitioning
25 the input query sequence into a plurality of overlapping windows.

26 8. The computer implemented process of claim 1 wherein searching the index
27 using the input query sequence further comprises determining the nearest neighbor of the
28 input query sequence.

29 9. The computer implemented process of claim 1 wherein searching the index
30 using the input query sequence further comprises:

31 partitioning the input query sequence into a plurality of windows; and
32 comparing one of the plurality of windows of the input query sequence with a
33 node of the index.

34 10. In a computer having a processor and a storage device, the storage device
35 storing a plurality of database sequences in computer-readable form, each database
36 sequence representing a biological sequence of nucleotides, a computer-implemented
37 process for selecting a database sequence in response to an input query sequence
38 representing a biological sequence of nucleotides, the process comprising:

39 generating a plurality of window vectors representing the database sequences;
40 generating a tree-structured index of the database sequences using the window
41 vectors; and
42 using the input query sequence to search the tree-structured index to identify
43 database sequences similar to the query sequence.

44 11. The computer-implemented process of claim 10 wherein generating a
45 plurality of window vectors further comprises:

46 partitioning each database sequence into a plurality of overlapping windows,
47 each window comprising a plurality of k-tuples; and
48 mapping the occurrence of each k-tuple in each window into a window vector.

49 12. The computer-implemented process of claim 11 wherein each of the
50 plurality of windows has a fixed length W and the overlap of the plurality of windows being
51 determined by a parameter Δ .

52 13. The computer implemented process of claim 10 further comprising
53 partitioning the input query sequence into a plurality of windows.

54 14. The computer implemented process of claim 10 wherein the tree-structured
55 index of the database sequences is generated using Tree-Structured Vector Quantization.

1 15. The computer implemented process of claim 10 wherein the tree-structured
2 index is a binary-tree index.

3 16. In a computer having a processor and storage, said storage comprising a
4 plurality of database sequences in computer-readable form, a computer-implemented
5 process for generating a tree-structured index representing the database sequences, the
6 process comprising:

7 partitioning each database sequence into a plurality of overlapping windows;
8 mapping each window into a window vector; and
9 generating the tree-structured index for the overlapping windows using the
10 window vectors.

11 17. The computer implemented process of claim 16 wherein each window has a
12 fixed length W and an offset Δ , each window beginning at position $j * \Delta$ and ending at
13 position $j * (\Delta + W)$, where $j = 0, 1, 2, 3....$

14 18. The computer implemented process of claim 16 wherein the tree-structured
15 index is generated using Tree-Structured Vector Quantization.

16 19. The computer implemented process of claim 16 wherein the tree-structured
17 index is a binary-tree index.

18 20. The computer implemented process of claim 16 further comprising:
19 in response to receiving a query sequence, searching the tree-structured index
20 using the query sequence to identify database sequences similar to the
21 query sequence.

22 21. In a computer-implemented system comprising a plurality of database
23 sequences, a system for generating a tree-structure index that represents the database
24 sequences, the system comprising:

25 a computer-readable medium for storing program instructions for execution by a
26 processor;
27 a window vector generating component, stored on the computer-readable
28 medium and executable by the processor, for mapping each database
29 sequence into a window vector; and
30 a tree-structure index generator, stored on the computer-readable medium and
31 executable by the processor, for generating a tree-structure index of the
32 window vectors.

33 22. The system of claim 21 wherein the k-tuple vector mapping component
34 further comprises:

35 a window partitioning component, stored on the computer-readable medium and
36 executable by the processor, for partitioning each database sequence into
37 a plurality of windows; and
38 a window vector mapping component, stored in the computer-readable medium
39 and executable by the processor for mapping each window into a window
40 vector.

41 23. A computer-implemented system for selecting database sequences from a
42 database comprising a plurality of sequences stored in electronic form, the system
43 comprising:

44 a computer processor for executing programmed instructions and for storing and
45 retrieving data;
46 an input device, coupled to the processor, for receiving a query sequence;
47 a program memory, coupled to the processor, for storing program instructions
48 for execution by the processor;
49 a tree-structure index generator, executable by the processor in accordance with
50 program instructions in the program memory, for generating a tree-
51 structure index that represents the database sequences;
52 a query sequence partitioning component, executable by the processor in
53 accordance with program instructions in the program memory, for
54 partitioning a query sequence into a plurality of windows; and

55 a retrieval component, executable by the processor in accordance with program
56 instructions in the program memory, for retrieving database sequences
57 from the database that are similar to the query sequence.

58 24. A computer readable memory, for use with a processor and a plurality of
59 database sequences, the memory storing a computer program executable by the processor
60 for configuring and controlling the processor to generate a tree-structured index that
61 represents the database sequences, by performing the steps of:

62 partitioning each database sequence into a plurality of overlapping windows;
63 for each window, generating a window vector representing the window; and
64 generating a tree-structured index of the window vectors.

65 25. A computer readable memory, for use in a computer system comprising a
66 processor and a database storing a plurality of database sequences in computer-readable
67 form, and for selecting database sequences from the database in response to a query
68 sequence, the computer readable memory storing:

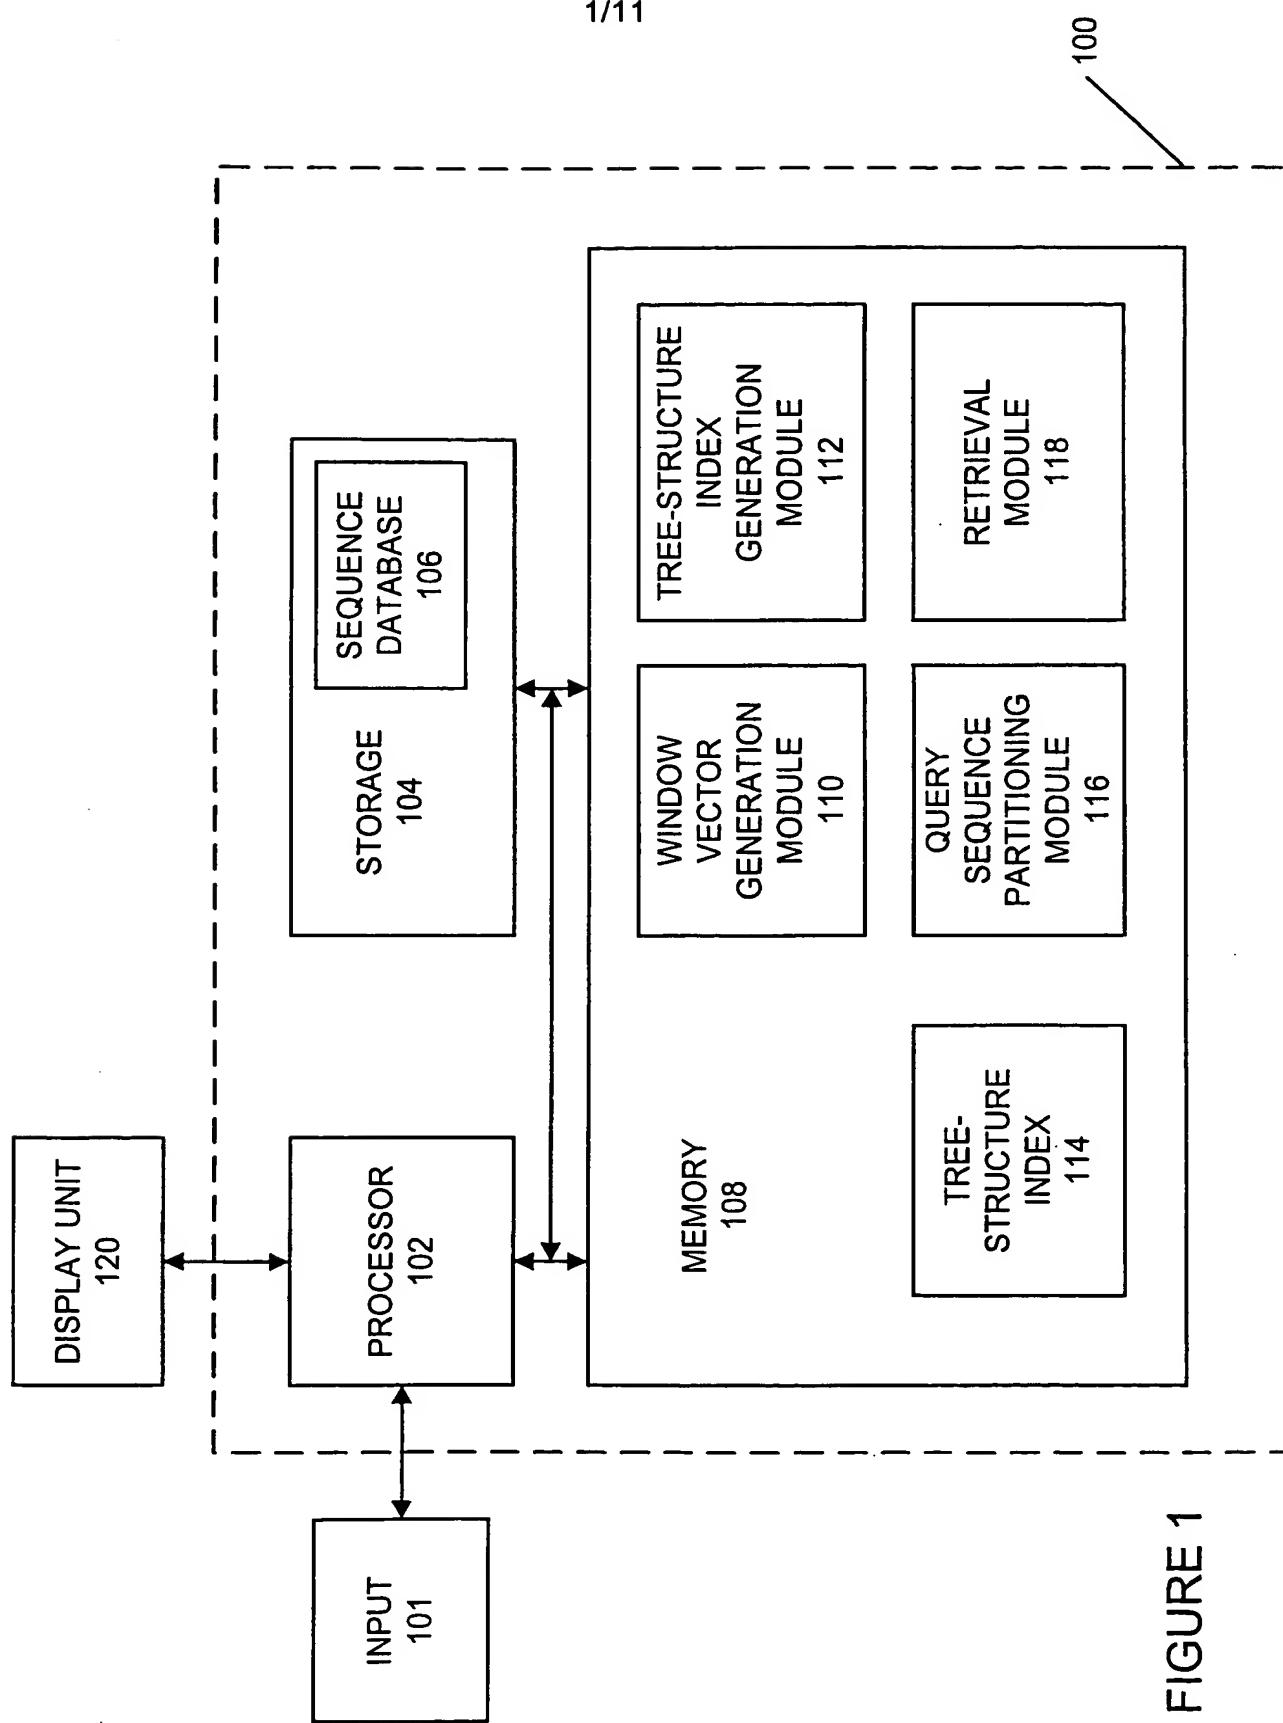
69 a tree-structured index generator, executable by the processor in accordance with
70 program instructions in the memory, for generating a tree-structured
71 index that represents the database sequences;

72 a query sequence partitioning component, executable by the processor in
73 accordance with program instructions in the memory, for partitioning a
74 query sequence into a plurality of windows; and

75 a retrieval component, executable by the processor in accordance with program
76 instructions in the memory, for retrieving database sequences from the
77 database that are similar to the query sequence.

78

1/11

**FIGURE 1**

2/11

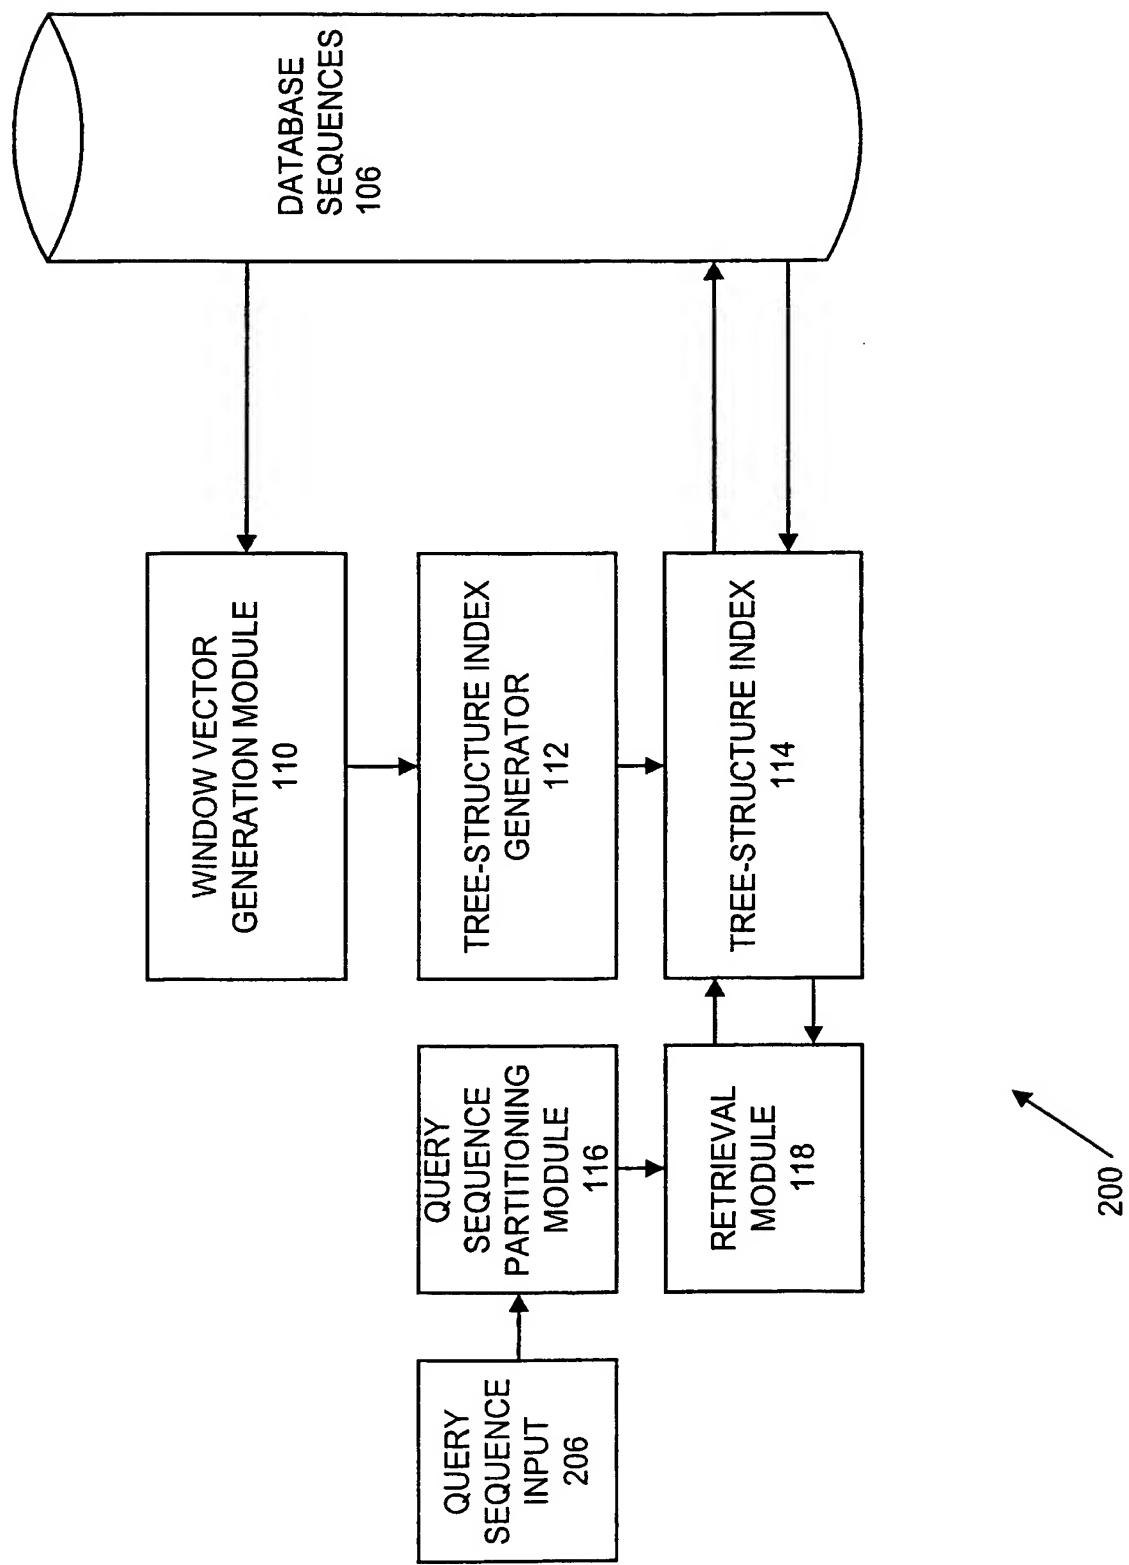


FIGURE 2

3/11

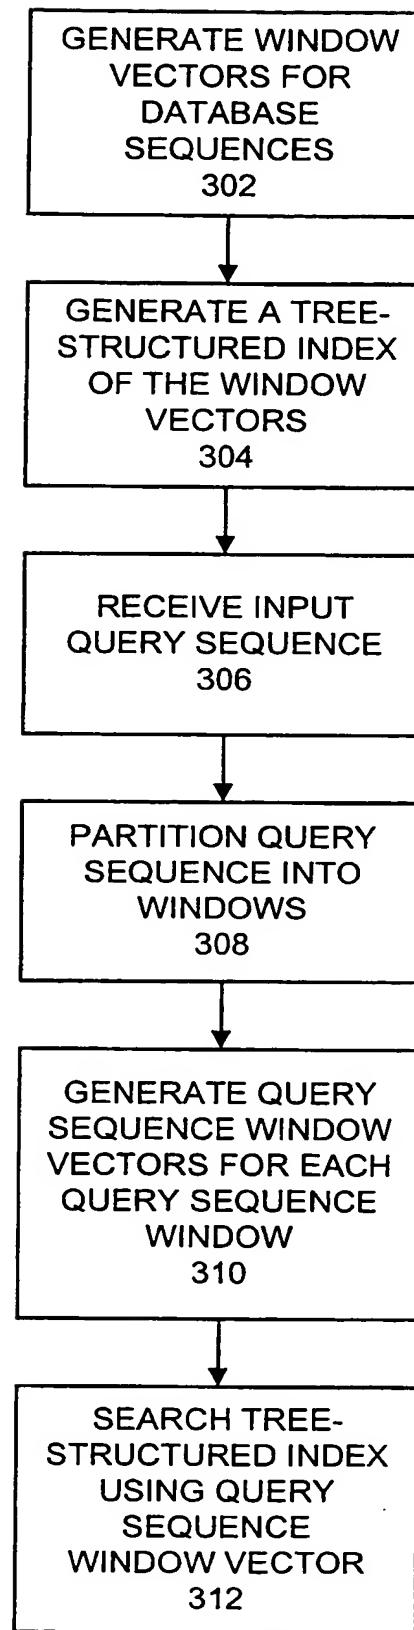


FIGURE 3

4/11

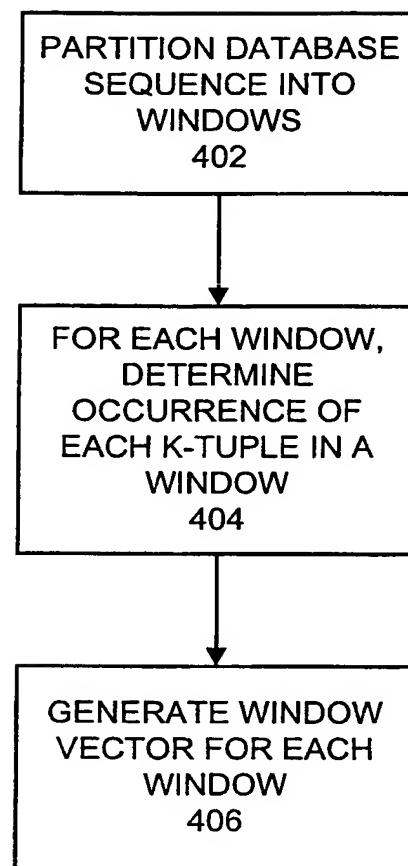
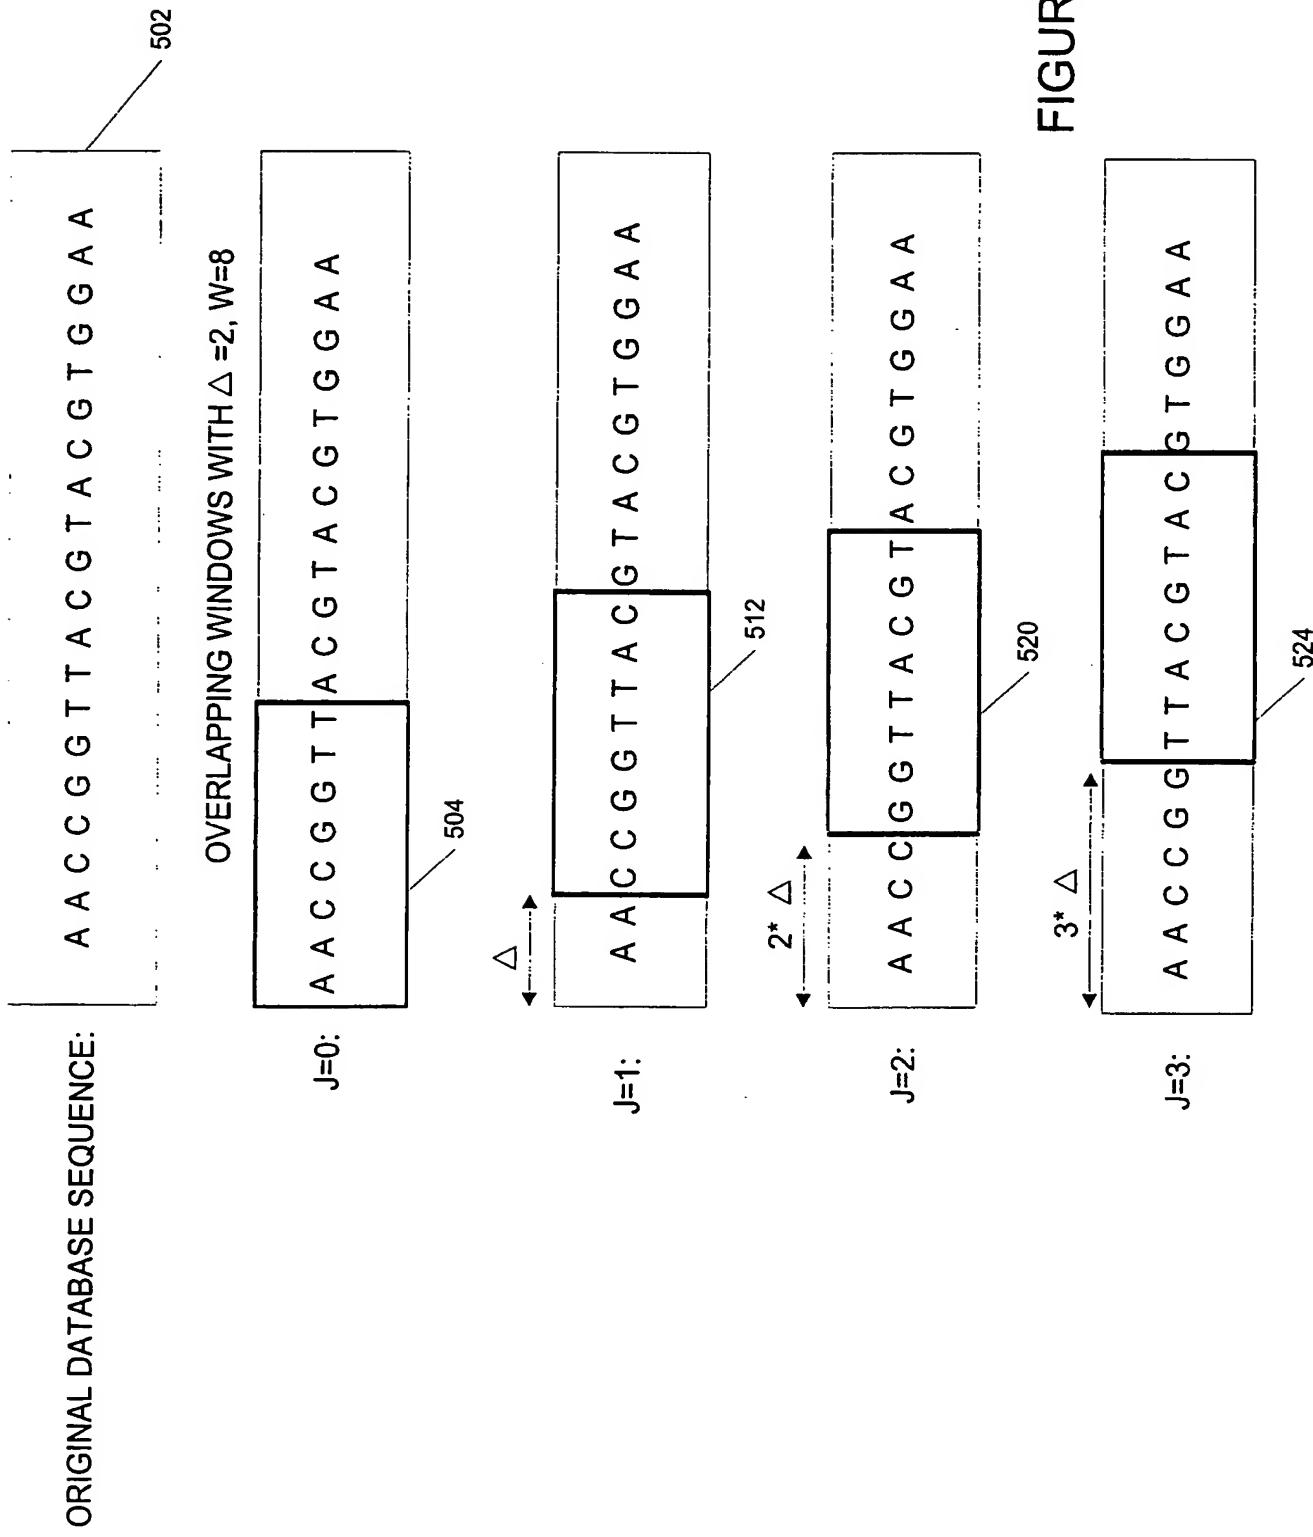


FIGURE 4

5/11



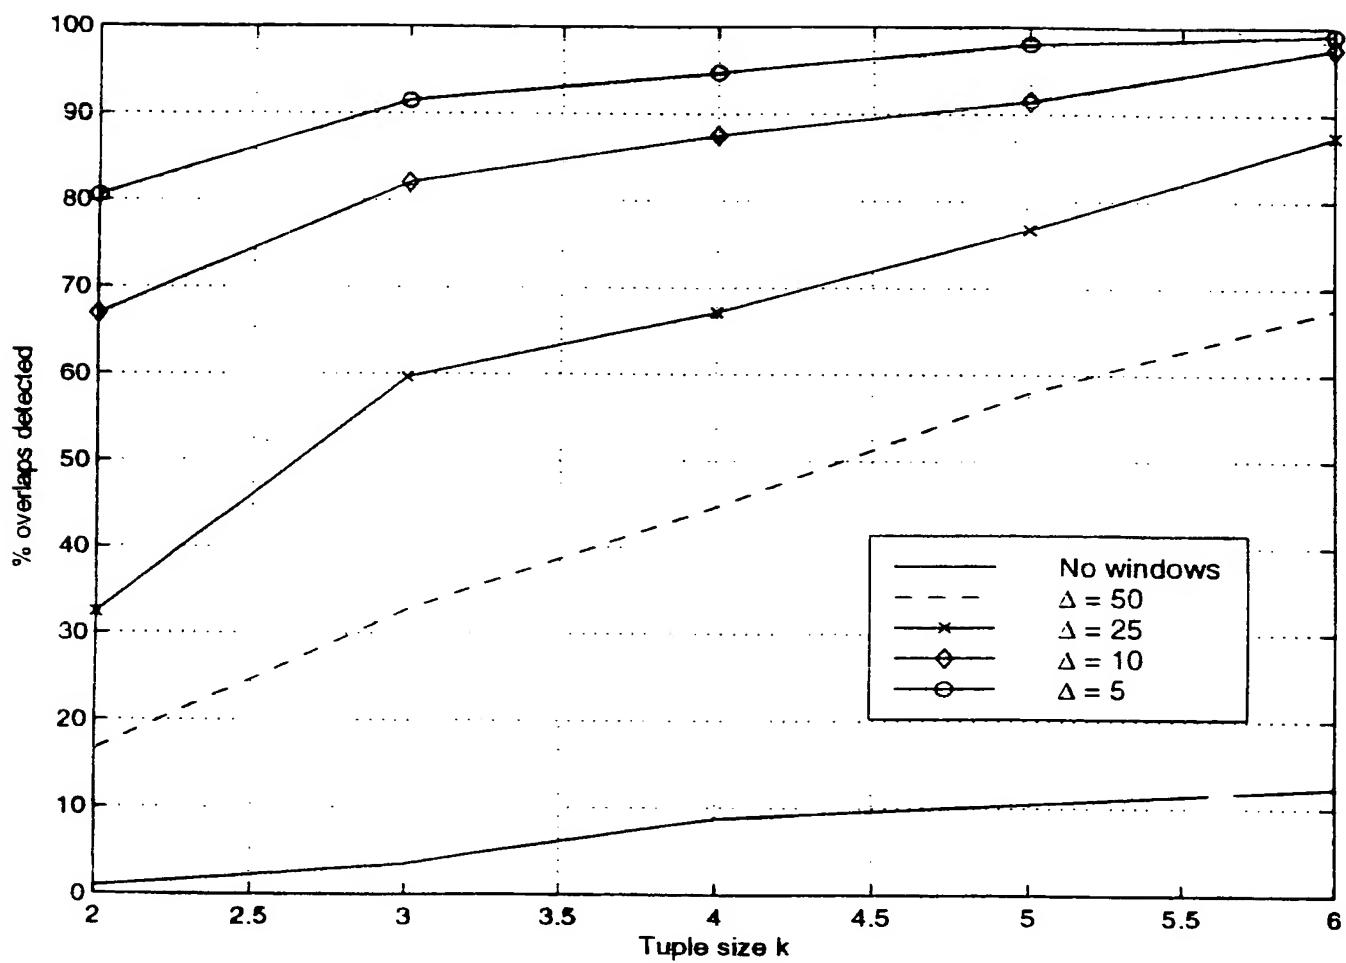


FIGURE 6

7/11

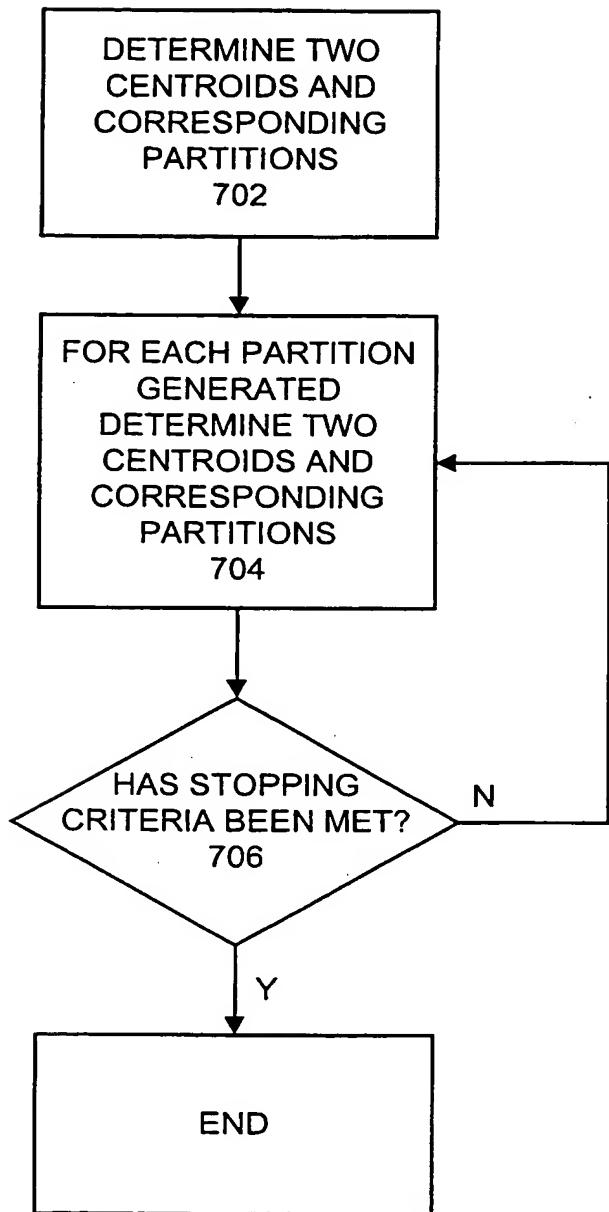


FIGURE 7

8/11

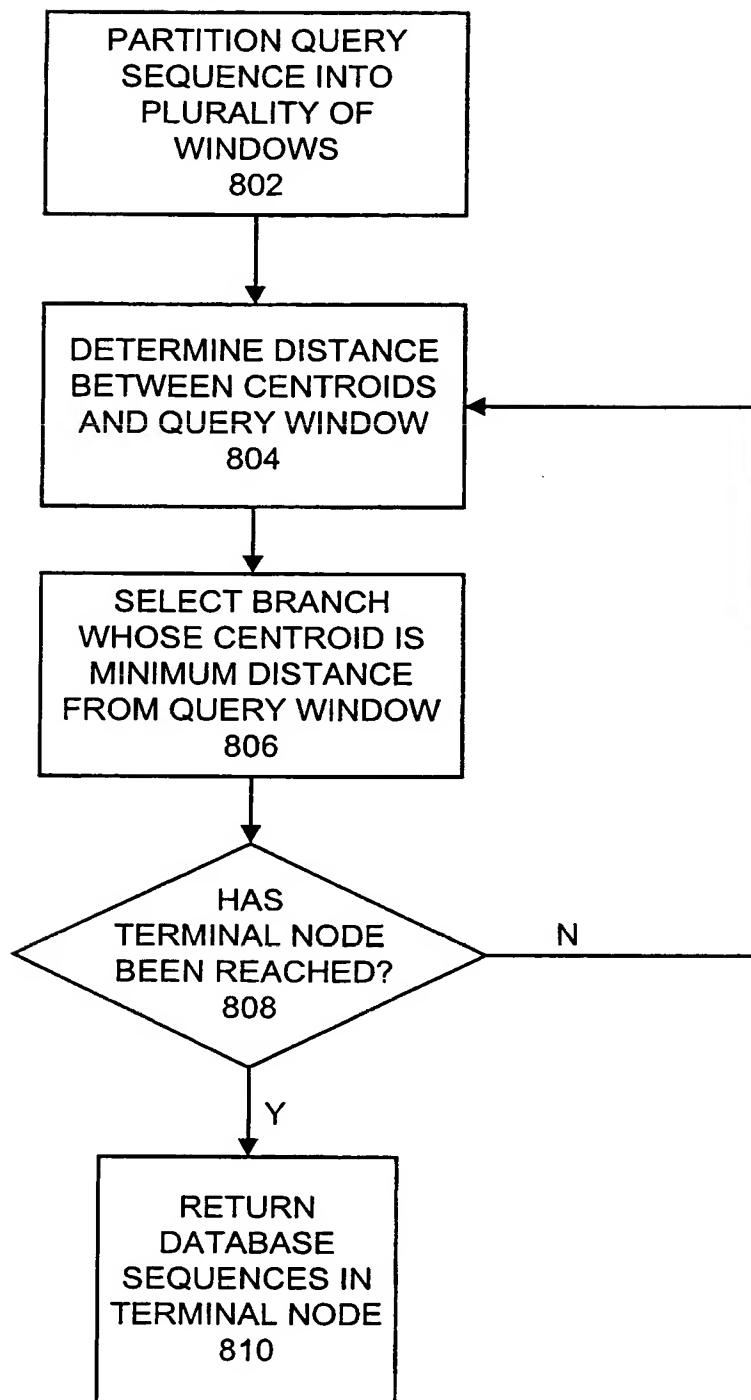


FIGURE 8

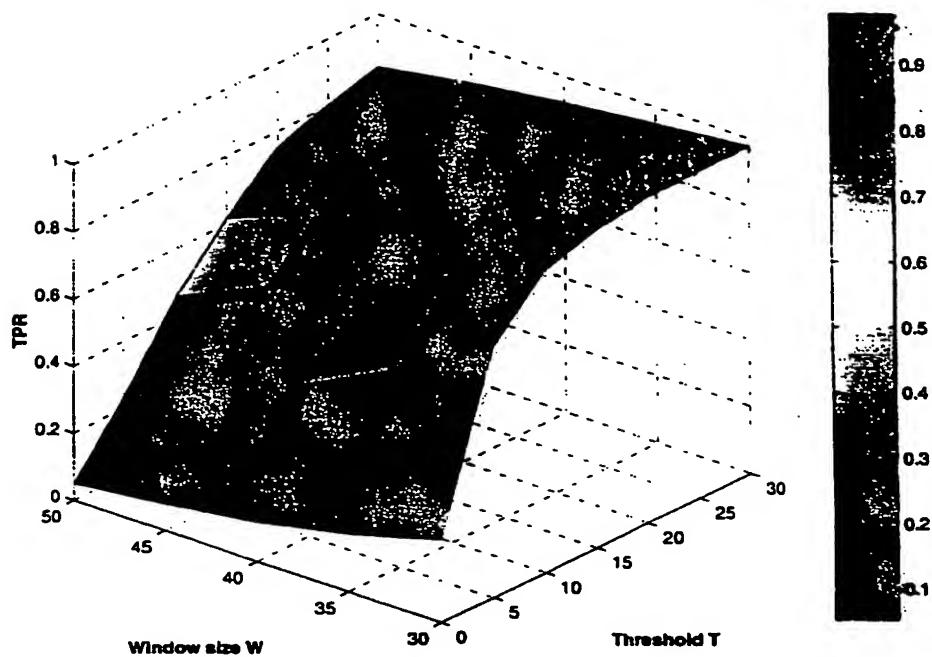


Figure 9

Figure 1: The TPR as a function of the window size W and the distance threshold T . Data has an error rate of 5 %.

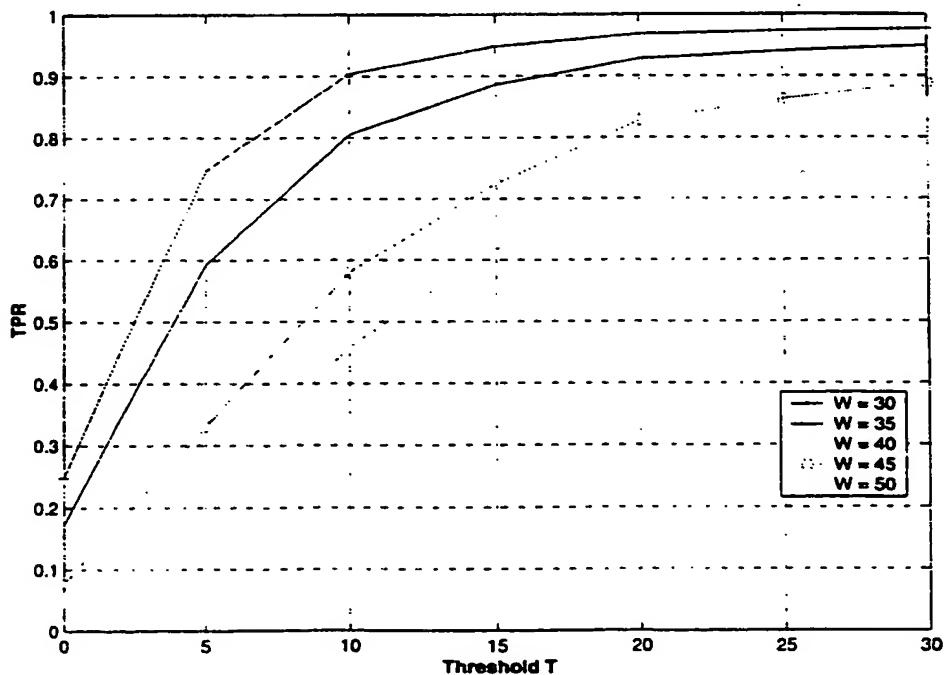


Figure 10

Figure 2: The ordinate is the TPR and the abscissa is the distance threshold T . Each graph corresponds to a different window size. Data has an error rate of 5 %.

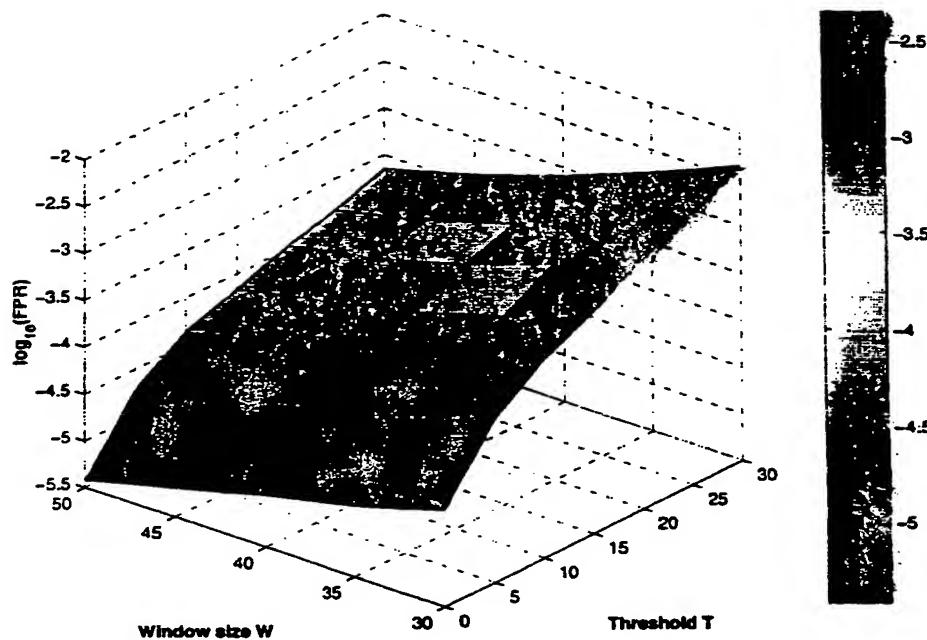


Figure 11

Figure 3: The log base 10 of the FPR as a function of the window size W and the distance threshold T . Data has an error rate of 5 %

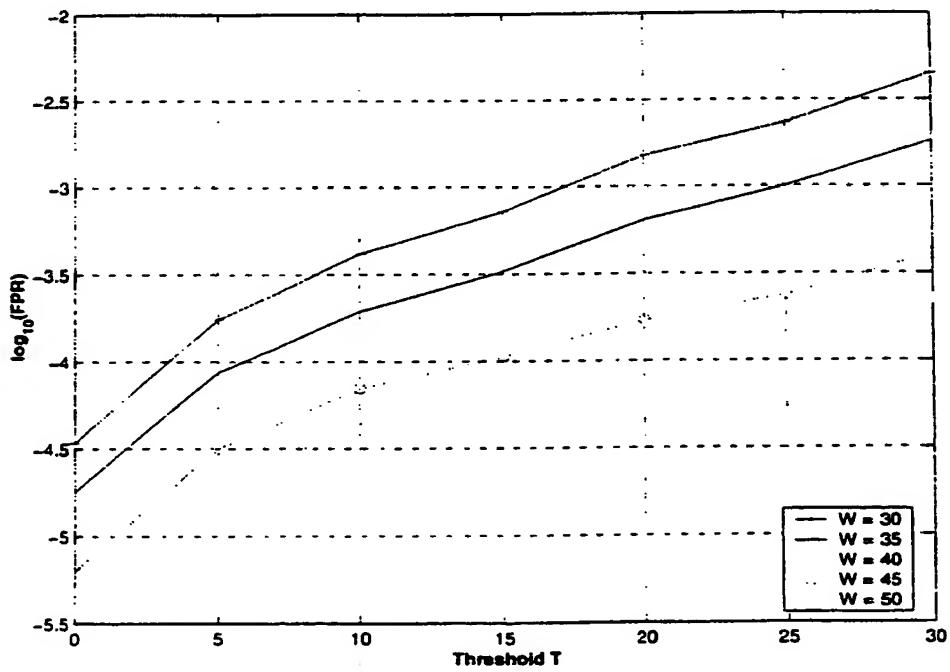


FIGURE 12

Figure 4: The ordinate is the log base 10 of the FPR and the abscissa is the distance threshold T . Each graph corresponds to a different window size. Data has an error rate of 5 %.

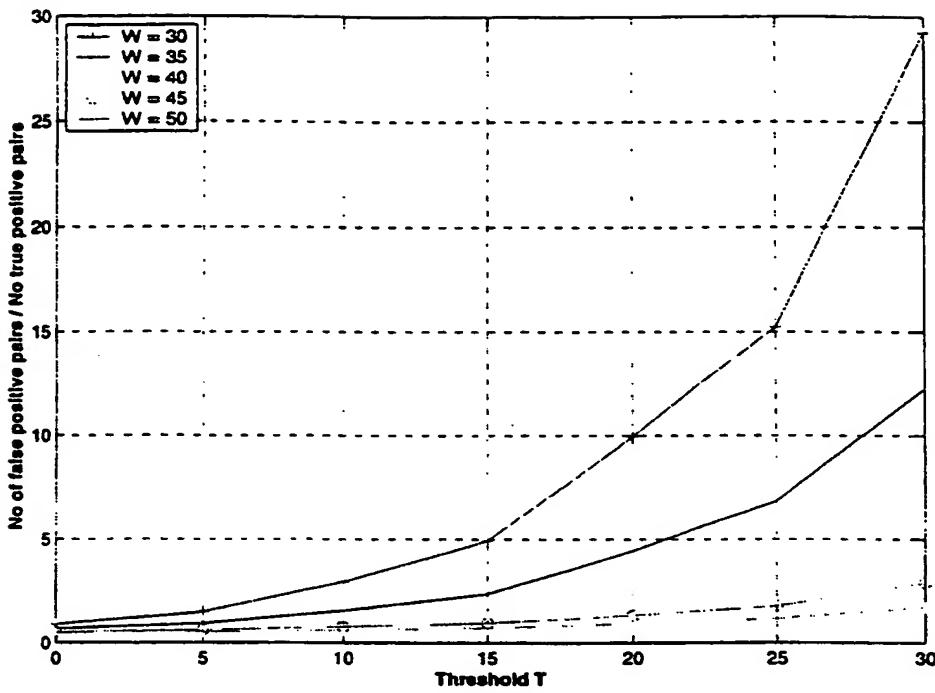


FIGURE 13

Figure 5: The ratio of the number of false positive pairs to the number of true positive pairs. Data has an error rate of 5 %.

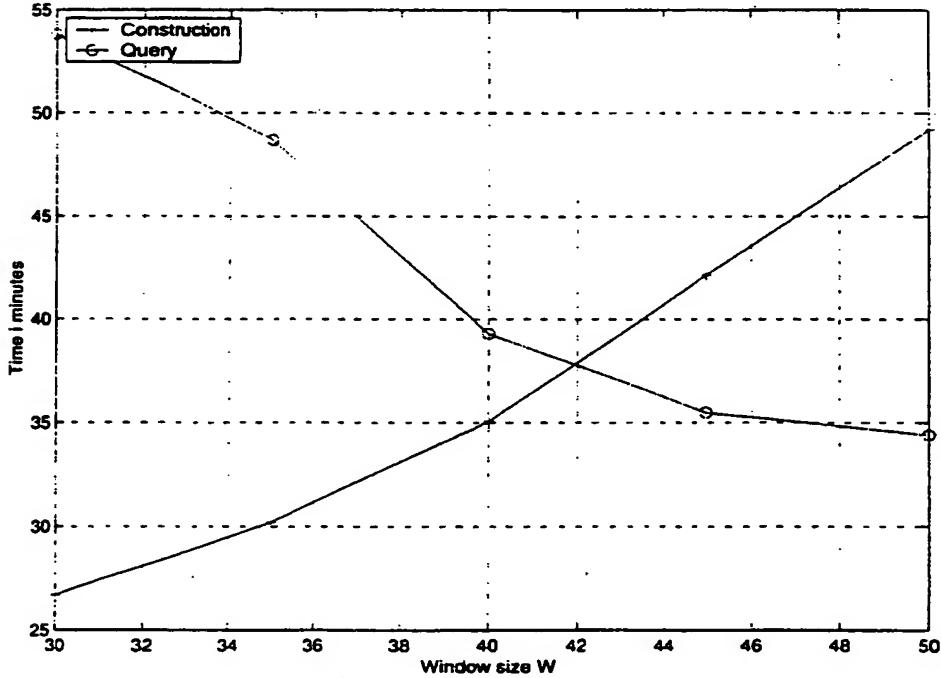


FIGURE 14

Figure 6: The ordinate is the computation time in minutes. The abscissa is the window size.

INTERNATIONAL SEARCH REPORT

Internat'l Application No
PCT/US 00/34413

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F G06C C12Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ, WPI Data, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 577 249 A (CALIFANO ANDREA) 19 November 1996 (1996-11-19) column 3, line 12 -column 4, line 63 ---	1,10,16, 21,23-25
A	US 5 701 256 A (CHANG WILLIAM I-WEI ET AL) 23 December 1997 (1997-12-23) abstract column 4, line 42 -column 5, line 47 ---	1,10,16, 21,23-25
A	EP 0 583 559 A (IBM) 23 February 1994 (1994-02-23) column 3, line 29 -column 4, line 37 ---	1,10,16, 21,23-25

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

° Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

23 April 2001

Date of mailing of the international search report

02/05/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+31-70) 340-3016

Authorized officer

Deane, E

INTERNATIONAL SEARCH REPORT

Internat'l Application No
PCT/US 00/34413

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	PATENT ABSTRACTS OF JAPAN vol. 1996, no. 05, 31 May 1996 (1996-05-31) & JP 08 022475 A (MITSUBISHI ELECTRIC CORP), 23 January 1996 (1996-01-23) abstract -----	1,10,16, 21,23-25

INTERNATIONAL SEARCH REPORT

Information on patent family members

Intern	al Application No
PCT/US 00/34413	

Patent document cited in search report	Publication date	Patent family member(s)			Publication date
US 5577249 A	19-11-1996	EP	0583559 A		23-02-1994
		JP	2673091 B		05-11-1997
		JP	6098770 A		12-04-1994
US 5701256 A	23-12-1997	CA	2177764 A		01-12-1996
EP 0583559 A	23-02-1994	JP	2673091 B		05-11-1997
		JP	6098770 A		12-04-1994
		US	5577249 A		19-11-1996
JP 08022475 A	23-01-1996		NONE		

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

CORRECTED VERSION

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
5 July 2001 (05.07.2001)

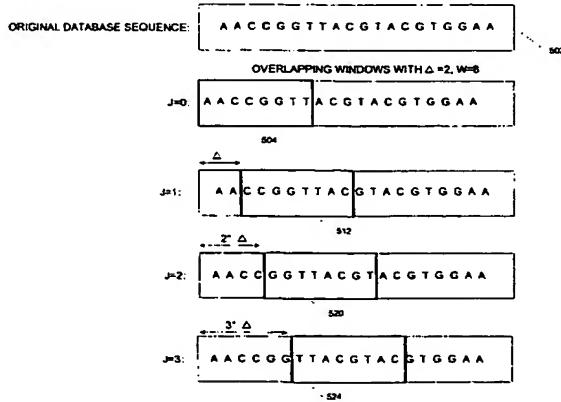
PCT

(10) International Publication Number
WO 01/48637 A1

- (51) International Patent Classification⁷: **G06F 17/30**
- (21) International Application Number: PCT/US00/34413
- (22) International Filing Date:
18 December 2000 (18.12.2000)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
09/474,929 29 December 1999 (29.12.1999) US
- (71) Applicants: INCYTE GENOMICS, INC. [US/US]: 3160 Porter Drive, Palo Alto, CA 94304 (US). STANFORD UNIVERSITY [US/US]: Office of Technology Licensing, Suite 350, 900 Welch Road, Palo Alto, CA 94304-1850 (US).
- (72) Inventors: WALKER, Michael, G.: 1050 Borregas Ave., #80, Sunnyvale, CA 94089 (US). WANG, James, Z.: 672
- (74) Agents: TRINIDAD, Arriola-Kern et al.: Fenwick & West LLP, Two Palo Alto Square, Palo Alto, CA 94306 (US).
- (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CR, CU, CZ, DE, DK, DM, DZ, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, TZ, UA, UG, UZ, VN, YU, ZA, ZW.
- (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: SEQUENCE DATABASE SEARCH WITH SEQUENCE SEARCH TREES



WO 01/48637 A1

(57) Abstract: A method and system for generating and searching a tree-structured index of window vectors that represent database sequences comprise a window vector generation module, a tree-structured index generation module, a query sequence partitioning module, and a retrieval component. The window vector generation module partitions a database sequence into a plurality of overlapping windows. Each window has a fixed length W comprising a fixed number of nucleotides, and the offset among windows is determined by a parameter Δ . The window vector generation module then maps each database sequence window into a window vector. The database sequence window vector indicates the frequency of appearance of each k-tuple in the corresponding database sequence window. The tree-structured index generation module then generates a tree-structured index using the database sequence window vectors. The query sequence partitioning module partitions a query sequence into a plurality of windows and maps each query sequence window into a query sequence window vector. Each query sequence window vector is then compared against the tree-structured index to locate the database sequences that are similar to the query sequence. The list database sequences that are similar to the sequence is then returned as the result of the search.



Published:

— *with international search report*

(15) Information about Correction:

see PCT Gazette No. 22/2002 of 30 May 2002, Section II

(48) Date of publication of this corrected version:

30 May 2002

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

SEQUENCE DATABASE SEARCH WITH SEQUENCE SEARCH TREES

5

GOVERNMENT SUPPORT

Work for the following disclosed invention was supported in part by the National
10 Science Foundation under Grant No. _____. The Government may have certain rights in
the invention.

BACKGROUND

A. Field of the Invention

The present invention relates generally to storage, indexing, and retrieval of
15 database sequences, and more particularly to a method and system for generating and
searching a tree-structured index of window vectors that represent database sequences.

B. Background of the Invention

Sequence-similarity finding programs identify sequences in DNA and
protein databases that are similar to a query sequence. Because of the recent explosion in
20 the amount of DNA sequence information available in public and private databases as a
result of the human genome project and other large-scale DNA sequencing efforts, such
sequence-similarity finding programs have become increasingly important in modern
biology.

Generally, there are two classes of sequence-similarity searching programs:
25 global comparison methods (e.g. Needleman and Wunsch), and local comparison methods,
e.g., the FASTA method (Pearson and Lipman), and the BLAST method (Altschul). Global
comparison methods have a high degree of accuracy but are extremely slow. Local
comparison methods such as FASTA and BLAST identify candidate similar sequences
based on shared k-tuples and therefore are faster than global methods. However, local
30 comparison methods are less accurate, i.e. they provide results which are not similar.
Moreover, with both the global and local comparison methods, the computational

complexity of such methods usually increases linearly with the size of the number of sequences to be searched. This is due in part to the fact that most prior art methods search all, or at least a very large part, of the sequence database. To improve the searching time, several prior art methods have created sequence databases using clusters of sequences, 5 including tree-structured indexes. However, many of these methods use conventional sequence alignment methods such as BLAST and FASTA to determine pairwise distances between the sequences. Thus, such cluster and tree-structured methods are also limited by the speed of the alignment methods described above.

10 Accordingly, what is needed is a system and method for organizing and searching database sequences that is fast and efficient, and at the same time provides a high degree of accuracy, that is, one that identifies sequences similar to a query sequence.

SUMMARY OF THE INVENTION

15 The present invention overcomes the limitations of conventional sequence-similarity-searching programs by using window vectors that represent database sequences in a sequence storage and retrieval system. The window vectors associated with the database sequences are organized into a tree-structured index for faster and more efficient searching of the database sequences. A query sequence is used to search the tree-structured index for database sequences that are similar to the query sequence.

20 In one embodiment, each database sequence is partitioned into a plurality of overlapping windows or fragments of fixed length. Each database sequence window has a fixed length W and the degree of offset between windows is determined by a parameter Δ . Each database sequence window comprises a subsequence of elements from the database sequence beginning at position $j*\Delta$, from the start of the database sequence, and ending at position $j*\Delta+W$, where $j = 0, 1, 2, \dots, (L-W)/\Delta$. In other words, each database sequence 25 window has a fixed length W which is advanced down the length of the original database sequence every Δ elements.

For each database sequence window, a database sequence window vector is computed. A database sequence window vector represents the occurrence of each k-tuple in the database sequence window. In one embodiment, the database sequence refers to a DNA 30 or protein sequence, and the occurrence of each k-tuple in the DNA or protein sequence window is represented by a vector of length 4^k . Each position in the vector represents a unique k-length sequence (i.e. k-tuple). If a k-tuple occurs more than one time in a database

sequence window, either the number of times that the k-tuple occurs may be recorded in the corresponding position in the vector or the value "1" may be recorded in the corresponding position in the vector to indicate that the k-tuple occurs at least once in the database sequence window. If a k-tuple does not occur in a database sequence, a zero in the 5 corresponding position in the vector may be used to indicate that the k-tuple does not occur in the sequence. In one embodiment, database sequence window vectors are stored in a tree-structured index to reduce searching time.

Database sequences are retrieved using a query sequence. In one embodiment, a query sequence is partitioned into a plurality of windows. For each query sequence 10 window, a query sequence window vector is computed. A query sequence window vector represents the occurrence of each k-tuple in the query sequence window. Each query sequence window vector is compared against the tree-structured index of the database sequence window vectors to locate the nearest neighbors of the query sequence, i.e. database sequences similar to the query sequence. In other words, for each query sequence 15 window vector, the tree-structured index is traversed from a root node of the tree to a terminal node which contains the nearest neighbor for that query sequence window vector. In one embodiment, the list of sequences with at least one significant window hit is returned.

In one embodiment, the present invention is used to generate and search a tree-20 structured index of window vectors that represent biological database sequences. Each database sequence may represent a DNA sequence comprising a fixed number of nucleotides. The DNA database sequence is then partitioned into a plurality of overlapping windows. Each DNA database sequence window has a fixed length W comprising a fixed 25 number of nucleotides, and the degree of offset among windows is determined by a parameter Δ . Each DNA database sequence window is then mapped into a database sequence window vector. The DNA database sequence window vector indicates the frequency of appearance of each k-tuple in the corresponding DNA database sequence window. A tree-structured index is then generated using the DNA database sequence window vectors. To search the tree-structured index, a query sequence (e.g. a DNA 30 sequence) is partitioned into a plurality of windows. Each query sequence window is then mapped into a query sequence window vector. Each query sequence window vector is then compared against the tree-structured index to locate the DNA database sequence window

vectors which are closest to the DNA query sequence window vector. The list of DNA database sequences that is similar to the DNA query sequence is then returned.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of a system in accordance with the present invention.

5 Fig. 2 is a functional block diagram of the present invention.

Fig. 3 is a flow chart of the steps of a preferred embodiment of the present invention.

Fig. 4 is a flow chart of the steps of a preferred embodiment for generating window vectors in accordance with the present invention.

10 Fig. 5 is an illustration of a preferred embodiment for partitioning a database sequence in accordance with the present invention.

Fig. 6 is a graph showing the results of computations performed to monitor the sensitivity of the present invention as a function of the offset size and the tuple size.

Fig. 7 is a flow chart of the steps of a preferred embodiment for generating a tree-structured index of the database sequences in accordance with the present invention.

15 Fig. 8 is a flow chart of the steps of a preferred embodiment for partitioning and using a query sequence to search a tree-structured index in accordance with the present invention.

Fig. 9 is a graph showing the true positive rate as a function of the window size.

Fig. 10 is a graph showing the true positive rate as a function of window size.

20 Fig. 11 is a graph showing the false positive rate as a function of the window size and threshold.

Fig. 12 is a graph showing the false positive rate as a function of the window size and threshold.

25 Fig. 13 is a graph showing the ratio of the number of false positive pairs to the number of true positive pairs.

Fig. 14 is a graph showing the time to generate a tree-structured index and the time to search it as a function of window size.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

A. System Overview

30 Referring now to Figure 1, there is shown a block diagram of software and hardware components for implementing one embodiment of the present invention 100. Processor 102

is a conventional engineering workstation or other computer processor and may also be coupled to other processors accessible over conventional communications channels or buses (not shown). Processor 102 is conventionally coupled to storage 104 which may be a magnetic disk storage, a CD storage unit, or other conventional computer data storage unit.

5 Storage 104 may also be coupled to other storage units accessible over conventional communications channels or buses (not shown). Storage 104 includes sequence database 106. Sequence database 106 contains a plurality of sequences, also referred to herein as "database sequences", to be searched. In a preferred embodiment, the database sequences represent biological sequences, such as DNA or protein sequences, that are represented in electronic or computer readable form. The present invention however is not limited to DNA

10 sequences, but may be used with other types of data.

Processor 102 is also conventionally coupled to memory 108 which is a random access memory (RAM) unit or other conventional computer memory. Items in memory 108 may alternatively be stored in storage 104 and accessed by processor 102 when required.

15 Memory 108 may comprise various software products including window vector generation 110, tree-structure index generation 112, tree structure index 114, query sequence partitioning 116, and retrieval 118 program components or modules as discussed below. Input 101 comprises conventional input devices such as a keyboard, mouse, track ball, or touchscreen. A conventional display unit 120 may also be conventionally coupled to

20 processor 102.

In a preferred embodiment, the present invention may be implemented on any platform, operating system, and user interface of sufficient power and flexibility, such as: XWindows/MOTIF; Sun/OS SunView; Microsoft Windows, VAX/VMS, and the like, and may be operated in a distributed data / distributed processor environment, if desired. In a preferred embodiment, the invention is capable of handling a very large data corpus, e.g., containing over 10 million sequences.

Referring now to Figures 2 and 3, there are shown a functional model of the present invention for generating window vectors and identifying database sequences similar to an input query sequence and a flow chart of the steps of a preferred use of the present invention. In this embodiment, the window vector generation module 110 generates 302 a plurality of window vectors associated with each database sequence. Preferably, the window vector generation module 110 partitions each database sequence into a plurality of windows and then encodes each window into a window vector. Preferred embodiments of

the functional operation of the window vector generation module 110 are described in more detail below with reference to Figure 4. The tree-structure index generation module 112 then generates 304 a tree-structured index 114 of the window vectors, preferably using Tree-Structured Vector Quantization (TSVQ). A preferred embodiment of the functional 5 operation of the tree-structure index generation module 112 is described below in more detail with reference to Figure 7. A query sequence 206 is input 306 into the system, and the query sequence partitioning module 116 partitions 308 the query sequence into a plurality of windows. In a preferred embodiment, the query sequence partitioning module 116 then maps or encodes 310 each query sequence window into a query sequence window 10 vector. A preferred embodiment of the functional operation of the query sequence partitioning module 116 is described below in more detail with reference to Figure 8. Finally, the retrieval module 118 then uses the query sequence window vector to search 312 the tree-structured index 114 to identify nearest neighbor window vectors for the query sequence, i.e. database sequences which are similar to the query sequence. A preferred 15 embodiment of the functional operation of the retrieval module 118 is described below in more detail with reference to Figure 8.

Thus, the present invention provides vector quantization and clustering methods for organizing databases for rapid and efficient searching. The present invention is based upon a database sequence search tree where the database sequences are represented using 20 database sequence window vectors organized into a tree-structured index. Window vectors represent the occurrences of each k-tuple of interest in a database sequence, or a query sequence, as described in more detail below. A k-tuple, as used herein, is a fixed length ordered sequence of elements and is well known to one skilled in the art of sequence-similarity searching. The use of window vectors advantageously increases the sensitivity of 25 the present invention as described in more detail below. The database sequence window vectors are then used to generate a tree-structured index representing the database sequences. The tree-structured index decreases the searching time of the present invention by allowing the present invention to search only a portion of the database sequences to find 30 similar sequences. Thus, whereas the searching time required for prior art sequence-similarity searching methods increases linearly with the size of the database, the searching time required for the present invention increases logarithmically with the size of the database.

B. Window Vector Generation

Referring now to Figure 4, there is shown a high-level flowchart of a preferred embodiment of the steps for generating database sequence window vectors in accordance with the present invention. Each database sequence comprises a plurality of elements, L, and is partitioned 402 into a plurality of windows or fragments. Preferably, a “sliding window scheme” is used to partition each database sequence into overlapping windows or fragments. In other words, each database sequence window has a fixed length W and the degree of offset between windows is determined by a parameter Δ . In one embodiment, the length of the window W refers to the number of DNA nucleotides or protein amino acids in the database sequence. Each database sequence window consists of subsequences of the original database sequence beginning at position $j*\Delta$, in the original database sequence, and ending at position $j*\Delta + W$, where $j = 0, 1 \dots (L-W)/\Delta$.

Referring now to Figure 5, there is shown an example of partitioning a database sequence in accordance with the sliding window scheme. For example, assume an original database sequence 502 comprises a DNA sequence of 20 nucleotides ($L=20$) where each nucleotide is represented by one of the letters A,C,G, or T. The original database sequence 502 is partitioned into a plurality of windows using the parameters $W=8$ and $\Delta = 2$. Accordingly, for $j=0$, window 504 is created. For $j=1$, window 512 is created. For $j=2$, window 520 is created. For $j=3$, window 524 is created. The process is continued until all the database sequence windows for the database sequence have been created, that is, until $j=(L-W)/\Delta$ (until $j=6$ in this example). Thus, for the example in Figure 5, seven overlapping windows will be generated for the original database sequence 502.

As can be seen in Figure 5, the offset between each window is a multiple of Δ . For example, the beginning of window 512 is offset from window 504 by Δ resulting in an overlap of the nucleotides CCGGTT. The beginning of window 520 is offset from window 504 by $2*\Delta$ and from window 512 by Δ . Therefore, between windows 520 and 504 there is an overlap of nucleotides GGTT, and between windows 520 and 512 there is an overlap of nucleotides GGTAC. Thus, each window has a fixed length W which is advanced down the length, L, of the original database sequence every Δ elements. In other words, each window runs from $j*\Delta$ to $j*\Delta + W$, where $j = 0, 1 \dots (L-W)/\Delta$. The sliding window scheme is used to generate a plurality of overlapping windows until the end of the

DNA database sequence is reached. In a preferred embodiment, values for Δ , W, and L are in the range $0 \geq W \geq 25$; $5 \geq \Delta \geq 50$; and $20 \geq L \geq 1000$.

Partitioning the database sequences into overlapping windows advantageously increases the sensitivity of the searching algorithm. Referring now to 5 Figure 6, there is shown the results of computations performed to monitor the sensitivity of the present invention as a function of the offset size, Δ , and the tuple size, k. The ordinate of Figure 6 is the percentage of overlaps detected and the abscissa is the tuple size used in the encoding. Each graphed line corresponds to a different offset size Δ in the sliding window encoding scheme. For example, the solid line corresponds to encoding with no 10 windows, while the dashed line, $\Delta=50$, corresponds to the use of a single offset. Each graphed line in the figure shows the percentage of overlaps correctly detected by the algorithm as a function of the tuple size k. Only overlaps of size greater than or equal than 50 were considered. Thus, as illustrated in Figure 6, the use of the "sliding window" scheme increases the sensitivity of the present invention. Moreover, as the tuple size k 15 increases, the sensitivity of the present invention also increases.

Referring again to Figure 4, each database sequence is partitioned into a plurality of overlapping windows and each database sequence window is mapped into a window vector. One skilled in the art will realize that the order in which the database 20 sequences are partitioned into a plurality of windows and the order in which the database sequence windows are mapped into database sequence window vectors is not essential to the invention. In other words, a single database sequence may be partitioned into a plurality of windows and afterwards, each database sequence window may then be used to compute the database sequence window vector corresponding to each database sequence window. In another embodiment, the present invention may first partition all the database sequences 25 into a plurality of database sequence windows, and afterwards compute the corresponding database sequence window vectors. In yet another embodiment, the present invention may partition a database sequence into a first database sequence window and then compute the database sequence window vector for the first database sequence window. The present invention may then partition the database sequence into a second database sequence 30 window and compute the database sequence window vector for the second database sequence window. Thus, one skilled in the art will realize that many options are available for the sequence of steps for partitioning each database sequence into a plurality of

overlapping windows and computing the database sequence window vectors corresponding to each database sequence window.

Referring back to Figure 4, there is shown one embodiment of a method for generating database sequence window vectors for a database sequence. After each database sequence has been partitioned 402 into a plurality of overlapping windows, each database sequence window is then mapped into a window vector. The database sequence window vectors are generated by first determining for each window 404 the frequency of appearance of each k-tuple in a specified database sequence window. A database sequence window vector is then generated 406 for the specified database sequence window. The database sequence window vector indicates the frequency of each k-tuple in that corresponding database sequence window. For a sequence, a k-tuple is a fixed length ordered sequence of elements.

In general, the number of possible k-tuples of length k is N^k , where N is the number of unique identifiers from which a sequence can be created and k is the tuple length. For example, from an alphabet of four letters, as in DNA for example, the total number of possible k-tuples is 4^k . For example, in a sequence generated from the DNA alphabet consisting of {A, C, G, T}, there are 16 possible ordered k-tuples of length k=2: {AA, AC, AG, AT, CA, CC, CG, CT, GA, GC, GG, GT, TA, TC, TG, TT}. In the case of protein sequences, there are 20 letters from which a protein sequence can be created, giving 20^k possible k-tuples. Alternate amino-acid alphabets with fewer letters would yield smaller numbers of k-tuples for a fixed k.

Referring to a DNA sequence as an example, the occurrence of each k-tuple in a DNA sequence may be represented by a window vector of length 4^k . For example, the DNA sequence AACTT contains 2-tuples AA, AC, CT, and TT. These occurrences of 2-tuples may be represented by the 16-position window vector {1100000100000001}. This window vector corresponds to the ordered representation {AA, AC, AG, AT, CA, CC, CG, CT, GA, GC, GG, GT, TA, TC, TG, TT}. Thus, the 1 in the first position corresponds to the presence of the k-tuple AA, the 1 in the second position corresponds to the presence of the k-tuple AC, the 1 in the eighth position corresponds to the presence of the k-tuple CT, and the 1 in the sixteenth position corresponds to the presence of the k-tuple TT. If a particular k-tuple is not observed in a sequence, then a zero may be placed in the corresponding position in the window vector, as is the case for AG, AT, etc.... If a k-tuple occurs more than one time in a sequence, either the number of times that the k-tuple occurs

may be recorded in the corresponding position in the window vector or the value “1” may be recorded in the corresponding position in the window vector to indicate that the k-tuple occurs at least once in the sequence. The k-tuple size may be of any length, but for DNA and protein database sequences, k is preferably in the range of $2 \leq k \leq 10$. Thus, window
5 vectors provide fast and accurate identification of similar sequences with acceptable space requirements.

In one embodiment, a sparse window vector representation may be used to generate the window vectors. Sparse window vector representation allows a user to store only the non-zero entries by storing (index, value) pairs rather than storing the whole vector.
10 For example, consider the DNA sequence ACCGTTGACACCC and suppose that k=2, $\Delta=2$, and W=5. On a disk or larger memory space, all tuples and an index to the beginning of each window are stored. For example, for the above DNA sequence, the following tuple information is stored: AC,CC, CG,GT,TT,TG,GA,AC,CA,AC,CC,CC. In a smaller memory space, such as a RAM, the windows are stored as sparse window vectors such as
15 $\{(AC, 1), (CC,1), (CG,1), (GT,1)\}$. Accordingly, when sparse window vectors are used, sparse window vector computations should also be used. In other words, only non-zero entries are used for computations. If an entry is zero in both vectors, it is zero when both vectors are subtracted and hence, no arithmetic needs to be done on entries that are zero. Thus, sparse window vector representation advantageously decreases the amount of
20 memory required for storing the vector windows. Additionally, by using sparse window vector computations, the present invention provides for faster and more efficient computation.

C. Tree-Structure Index Generator

The tree-structured index or database sequence search tree (“SST”) organizes the window vectors that represent the database sequences to enable fast and efficient searching of the database sequences. In a preferred embodiment, a binary-tree structured index is used. Alternatively, the tree-structured index may use 3-way branching, 4-way branching, or an arbitrary number of branches at each branch node in the tree. The binary-
25 tree structured index provides for fast searching but requires the largest amount of space. Thus, the branching factor may be selected to best suit the constraints of the available computer hardware.

The tree-structured index is generated by recursively searching the database sequence window vectors for clusters that provide binary (or higher order) partitions. A variety of methods are available for finding such clusters and building the tree-structured index as would be known to one skilled in the art. In one embodiment, a Tree-Structured Vector Quantization (TSVQ) method using k-means clustering is used to generate the tree-structured index. Referring now to Figure 7, there is shown a high-level flowchart of a preferred embodiment of the steps for generating a tree-structured index of the database sequence window vectors in accordance with the present invention. First, two centroids and their corresponding partitions of the data are determined 702. In a preferred embodiment, this is done by choosing and designating two points, a_0 and b_0 , in the k-dimensional space as centroids. As used herein, a centroid represents a group of database sequence window vectors. The set of all database sequences closest to each centroid comprises a partition of the k-dimensional space. For example, in a 2-dimensional subspace, the database sequences are partitioned into two disjoint subsets, A and B, by assigning to set A, the sequences closer to a_0 , and assigning to set B all the other sequences. In other words, for each database sequence, the distance of each sequence from each centroid is calculated and the sequence is assigned to the partition corresponding to the nearest centroid, i.e., the centroid, a_0 and b_0 , which is a minimum distance from the sequence. A new centroid is then determined for each partition by calculating the sum of the distances of all the sequences in the partition from the centroid using the following equations:

$$a_1 = \sum_{v \in A} v / |A| \quad b_1 = \sum_{v \in B} v / |B|$$

where A is the set of vectors in Group A, B is the set of vectors in group B, v is a vector, and $\sum_{v \in A}$ is the sum of all vectors belonging to set A.

The steps for determining the two centroids are repeated until convergence, that is, until the change in the sum of the distances at each iteration is less than a small threshold, ϵ , or no vectors change partition. In one embodiment, the following equation is used to determine convergence:

$$||a_0 - a_1|| < \epsilon$$

where ϵ is preferably a small positive number such as $\epsilon = 10^{-10}$.

If there is no convergence, then a_0 is set to a_1 and b_0 is set to b_1 and step 702 is repeated until there is convergence. For each partition generated, two centroids and their corresponding partitions are determined 704 using the steps described above. These steps 5 are repeated 706 until a stopping criterion is met, that is until the number of training sequences in the partitions is less than a user-specified threshold, or the maximum number of terminal nodes is reached. Each of the terminal nodes is a partition of the k-dimensional space and contains a set of sequences that are nearest neighbors to the centroid of that node. In a preferred embodiment, 16-bit finite precision arithmetic is used to compute the 10 centroids. However, one skilled in the art will realize that floating point arithmetic may also be used to compute the centroids.

One skilled in the art will realize that the centroids may be computed using a sample of the windows in the database rather than computing centroids for all the window vectors. For example, in one embodiment, a sample of 500 sequences to estimate two 15 centroids may provide an improvement in the speed for generating a tree-structured index while still providing a relatively low cost in the error rate. Alternatively, one skilled in the art will realize that where higher accuracy is required, a larger sample of sequences for estimating two centroids is preferable. Thus, the use of sampling in generating a tree-structured index enables one to reduce the false negative rate by generating two trees, using 20 different samples, and taking the union of the database sequences in the terminal nodes in which the query sequence falls in the two trees. Both the generation and search of both trees may be done in parallel thereby increasing efficiency and speed. Additionally, in one embodiment, generation of the tree-structured index is completed the first time a leaf size is smaller than 100.

25

D. Query Sequence and Database Sequence Retrieval

A query sequence is used to search the tree-structured index to find the database sequences that represent the nearest neighbors of the query sequence. In one embodiment a query sequence refers to a biological sequence, such as a DNA or protein 30 sequence, for which one wishes to find similar sequences in the sequence database. Referring now to Figure 8, there is shown a high-level flowchart of a preferred embodiment of the steps for partitioning and using the query sequence to search the tree-structured index. Each query sequence is first partitioned 802 into a plurality of windows. Each window of

the query sequence has a fixed length W which corresponds to the length of the database sequence windows. In one embodiment, the query sequence may be partitioned into non-overlapping windows. In another embodiment, the query sequence may be partitioned into overlapping windows where the degree of offset among windows is determined by the 5 parameter Δ as discussed above with respect to window vector generation. Preferably, the query sequence is partitioned using either value for Δ : $\Delta = W$ or $\Delta = W/2$. Each query sequence window is then mapped into a window vector. A query sequence window vector represents the occurrence of each k-tuple in the query sequence window.

Each query sequence window vector is then used to search the tree-structured 10 index by choosing the node whose centroid is the minimum distance between the query sequence window vector and the centroid of that node. The list of database sequences with at least one significant window hit is returned. In other words, the database sequence window vectors in a leaf which are closer to query sequence window vectors than to a given threshold are returned. Additional processing such as alignment of the query sequence to 15 the database sequence with one of the standard alignment tools is also possible. In other words, one can use the Smith-Waterman algorithm to align windows to obtain a better score.

Referring again to Figure 8, the steps for searching a binary tree-structured 20 index are shown. A search of the tree-structured index begins at the root node of the tree. At each non-terminal node in the tree, each branch is represented in the tree by its centroid. The distance between each centroid and the query window is determined 204. In a preferred embodiment, the distance, d , between two windows, such as a database sequence window, W_1 , and a query window, W_2 , is determined by the following equation:

$$25 \quad d(W_1, W_2) = || X(W_1) - X(W_2) ||$$

where $X(W)$ refers to the window vector of the window.

The branch whose centroid is the minimum distance from the query sequence window is selected 206. Steps 204 and 206 are repeated through the tree, progressing from 30 the root note towards the terminal nodes, until a terminal node is reached 208. The database sequences in the selected terminal node are the nearest neighbors of the query sequence window. The database sequences representing the nearest neighbors of the query sequence window are then retrieved 210 from the sequence database. In other words, for each query

sequence, all database sequences which have a window in common with the query sequence are outputted as a resulting hit.

In another embodiment, the query sequence may be used to search only a portion of the tree-structured index, i.e. a subtree. For example, when searching the tree-structured index using the query sequence window, all query sequence windows may be used to search against a top portion of the tree-structured index. For example, all of the windows of a query sequence may be used to search only the first nine levels of the tree-structured index. Such an embodiment would generate $2^9 = 512$ groups of windows, where each group contains query sequence windows for the search completed in a different leaf of the top subtree. In such an embodiment, each group is searched against the subtree where it hits. With this embodiment, disk access occurs only when the data associated with each subtree is loaded into memory. One skilled in the art will realize that this embodiment may use an arbitrary number of levels for the subtree and that parallel searching of the tree-structured index may be utilized. For example, after the top subtree is processed, lower level subtrees may be processed using different processors. Thus, the present invention provides a faster and efficient method for searching the tree-structured index.

As the number of database sequences increases, it may become difficult to store all the window vectors in memory. In one embodiment, all vectors are stored on a disk and loaded into RAM as needed. In another embodiment, a caching strategy is used. Due to the in-order nature of the tree-structure generation, a FIFO cache may be used. In such embodiment, a refresh mechanism in which a hit in the cache also moves the vector to the front of the queue may also be used. Additionally, the window vector file should be reordered such that all vectors associated with a certain subtree are contiguous.

25

E. Performance Results

Tables 1, 2, and 3 below show the performance results of the present invention in comparison with the prior art method BLAST. In the computations, a 1.552KB piece of genomic DNA was fragmented several times using a Poisson process with average $\Lambda = 300\text{bp}$. From the pool of fragments, three sets of 30675, 61350, and 122700 sequences, respectively, were generated. The present invention and the BLAST2 prior art method were used to determine which of the database sequences overlapped. For the computation using the present invention, a window size of length 50 was used and the offset step $\Delta=5$. The

computation was repeated for several values of the tuple size k, i.e. k = 3, 4, 5, 6. Tables 1, 2, and 3 show the results of the present invention on each of these sets.

Table 1

K	Total	Tree Generation	Query	TPR	FNR	TNR	FPR
3	00:10:12	00:06:56	00:03:16	0.893	0.107	9.9998e-01	2.092e-05
4	00:11:37	00:07:24	00:04:13	0.932	0.068	9.9999e-01	1.069e-05
5	00:13:52	00:08:24	00:05:28	0.956	0.044	9.9999e-01	6.642e-06
6	00:23:04	00:13:04	00:10:00	0.963	0.037	9.9999e-01	5.284e-06

5 Table 1 shows the results of the present invention with 30,675 sequences. The time for generating the tree-structured index, for tree search, and the total time are displayed for k=3 to 6. The true positive rate (TPR), false negative rate (FNR), true negative rate (TNR), and false positive rate (FPR) are also displayed.

Table 2

K	Total	Tree Generation	Query	TPR	FNR	TNR	FPR
3	00:22:33	00:15:36	00:06:57	0.886	0.114	9.9999e-01	1.842e-05
4	00:26:11	00:17:37	00:08:34	0.926	0.074	9.9999e-01	9.805e-06
5	00:36:49	00:19:00	00:17:49	0.953	0.047	9.9999e-01	6.222e-06
6	00:45:53	00:26:37	00:19:16	0.960	0.040	9.9999e-01	5.018e-06

10 Table 2 shows the results of the present invention with 61,350 sequences. The time for generating the tree-structured index, for tree search, and the total time are displayed for k=3 to 6. The true positive rate (TPR), false negative rate (FNR), true negative rate (TNR), and false positive rate (FPR) are also displayed.

Table 3

K	Total	Tree Generation	Query	TPR	FNR	TNR	FPR
3	00:46:56	00:31:52	00:15:04	0.881	0.119	9.9999e-01	1.529e-05
4	00:50:50	00:32:50	00:18:00	0.924	0.076	9.9999e-01	8.404e-06

5	01:08:22	00:35:21	00:33:01	0.951	0.049	9.9999e-01	5.950e-06
6	01:34:50	00:52:21	00:42:29	0.957	0.043	1.0000e-00	4.793e-06

Table 3 shows the results of the present invention with 122,700 sequences. The time for generating the tree-structured index, for tree search, and the total time are displayed for k=3 to 6. The true positive rate (TPR), false negative rate (FNR), true negative rate (TNR), and false positive rate (FPR) are also displayed.

Searching the database sequences using the BLAST2 prior art method took 02: 03: 09, 07: 27: 47, 28: 10: 25 for 30,675, 61,350, and 122,700, respectively. Thus, the BLAST2 method yields an average time to bin each sequence of 0.2409, 0.4379, 0.8266 seconds. Thus, with the BLAST2 method there is a linear increase in time per sequence. Thus, as can be seen from the data, the search time required for the present invention is virtually independent of the number of database sequences and is for 122700 sequences 0.00074, 0.0088, 0.0161, and 0.0208 for k = 3, 4, 5, 6, respectively. Therefore, for 122,700 sequences the speed-up compared to BLAST2 varies between: 93 for k=3 to 39 for k=6. The search and construction time per sequence using the present invention is 0.0230, 0.0246, 0.0334, and 0.0464 for k = 3, 4, 5, 6, respectively. Therefore, for 122700, the speed-up compared to BLAST2 is 35 for k=3 to 17 for k=6.

F. Determination of Optimal Parameters

Experiments were performed to determine the optimal values for the threshold T, and the window size, W. The results described below were obtained with data with a 5% error rate. Referring now to Figure 9, there is shown a graph indicating the true positive rate (TPR) as a function of the window size, W, and the distance threshold T. Cross-sections of this surface, for fixed values of the window size W are presented in Figure 10. Both figures indicate that the TPR improves as the window size decreases and as the tolerance T increases. Moreover, Figure 9 shows that the TPR has an asymptotic value as T increases. The asymptotic value increases with a decreasing window size.

The FPR also increases with the tolerance T as indicated by Figures 11 and 12 in which the $\log_{10}(\text{FPR})$ is presented. On a logarithmic scale, the growth of the FPR is essentially linear, and the growth rate (the slope) is independent of the window size. Thus, the FPR grows exponentially as a function of T.

In a preferred embodiment, the window size W and the distance threshold T are chosen such that the TPR is large while the number of false positive pairs remain relatively small. Figure 13 indicates the ratio of the number of false positive pairs to the number of true positive pairs returned by the present invention. Thus, as can be seen from 5 Figure 13, for each window size, this ratio increases exponentially with the threshold T. In a preferred embodiment, the parameters W and T are chosen such that their ratio is bounded by 5. Among all possible values for these parameters, the values for which the true positive rate is maximized is desired. The values W=30 and T=15 has a ratio of about 2.5 and a TPR around .95, and the values W=35 and T=20 has a ratio less than 5 and a TPR around 10 .925.

Referring now to Figure 14, there is shown a graph indicating the time to generate a tree-structured index and the time to search it. The time for generating the tree-structured index increases with the window size. The trend for the search time is reversed. To explain this, one can see that the number of windows in the sequence database N is 15 about $N=n/\Delta$, where n is the total number of nucleotides in the database and Δ is the offset parameter. Therefore, N is independent of the window size. However, as the length of a window, W, increases, $X(W)$ has more non-zero entries because W contains more tuples. It follows that the time to compute the distance of $X(W)$ has more non-zero entries because W contains more tuples. Accordingly, the time to compute the distance of $X(W)$ to a centroid 20 increases with the length W, and which in turn increases the overall computation time in generating the tree-structured index. In a preferred embodiment, the number of query windows M per query is $M=m/M$ where m is the number of nucleotides in a query. Thus, as the window size increases, the number of query windows decreases so that the time per query also decreases.

We claim:

1 1. In a computer having a processor and storage, the storage comprising a
2 plurality of database sequences, a computer-implemented process for selecting database
3 sequences, the process comprising:

4 generating a plurality of window vectors representing the database sequences;

5 generating an index of the window vectors; and

6 in response to receiving an input query sequence, searching the index using the
7 input query sequence to select database sequences similar to the query
8 sequence.

9 2. The computer implemented process of claim 1 wherein generating a plurality
10 of window vectors further comprises:

11 partitioning each database sequence into a plurality of overlapping windows ;
12 and

13 generating a window vector for each window.

14 3. The computer implemented process of claim 2 wherein each database
15 sequence comprises a finite number of elements, L, and each window begins at position $j * \Delta$ and ends at position $j\Delta + W$, where $j = 0, 1 \dots (L-W)/\Delta$.

17 4. The computer implemented process of claim 1 wherein generating the index
18 further comprises:

19 searching the window vectors for clusters providing partitions.

20 5. The computer implemented process of claim 1 wherein the tree-index is a
21 tree-structured index.

22 6. The computer implemented process of claim 1 further comprising partitioning
23 the input query sequence into a plurality of non-overlapping windows.

24 7. The computer implemented process of claim 1 further comprising partitioning
25 the input query sequence into a plurality of overlapping windows.

26 8. The computer implemented process of claim 1 wherein searching the index
27 using the input query sequence further comprises determining the nearest neighbor of the
28 input query sequence.

29 9. The computer implemented process of claim 1 wherein searching the index
30 using the input query sequence further comprises:

31 partitioning the input query sequence into a plurality of windows; and
32 comparing one of the plurality of windows of the input query sequence with a
33 node of the index.

34 10. In a computer having a processor and a storage device, the storage device
35 storing a plurality of database sequences in computer-readable form, each database
36 sequence representing a biological sequence of nucleotides, a computer-implemented
37 process for selecting a database sequence in response to an input query sequence
38 representing a biological sequence of nucleotides, the process comprising:

39 generating a plurality of window vectors representing the database sequences;
40 generating a tree-structured index of the database sequences using the window
41 vectors; and
42 using the input query sequence to search the tree-structured index to identify
43 database sequences similar to the query sequence.

44 11. The computer-implemented process of claim 10 wherein generating a
45 plurality of window vectors further comprises:

46 partitioning each database sequence into a plurality of overlapping windows,
47 each window comprising a plurality of k-tuples; and
48 mapping the occurrence of each k-tuple in each window into a window vector.

49 12. The computer-implemented process of claim 11 wherein each of the
50 plurality of windows has a fixed length W and the overlap of the plurality of windows being
51 determined by a parameter Δ .

52 13. The computer implemented process of claim 10 further comprising
53 partitioning the input query sequence into a plurality of windows.

54 14. The computer implemented process of claim 10 wherein the tree-structured
55 index of the database sequences is generated using Tree-Structured Vector Quantization.

1 15. The computer implemented process of claim 10 wherein the tree-structured
2 index is a binary-tree index.

3 16. In a computer having a processor and storage, said storage comprising a
4 plurality of database sequences in computer-readable form, a computer-implemented
5 process for generating a tree-structured index representing the database sequences, the
6 process comprising:

7 partitioning each database sequence into a plurality of overlapping windows;
8 mapping each window into a window vector; and
9 generating the tree-structured index for the overlapping windows using the
10 window vectors.

11 17. The computer implemented process of claim 16 wherein each window has a
12 fixed length W and an offset Δ , each window beginning at position $j * \Delta$ and ending at
13 position $j * (\Delta + W)$, where $j = 0, 1, 2, 3, \dots$.

14 18. The computer implemented process of claim 16 wherein the tree-structured
15 index is generated using Tree-Structured Vector Quantization.

16 19. The computer implemented process of claim 16 wherein the tree-structured
17 index is a binary-tree index.

18 20. The computer implemented process of claim 16 further comprising:
19 in response to receiving a query sequence, searching the tree-structured index
20 using the query sequence to identify database sequences similar to the
21 query sequence.

22 21. In a computer-implemented system comprising a plurality of database
23 sequences, a system for generating a tree-structure index that represents the database
24 sequences, the system comprising:

25 a computer-readable medium for storing program instructions for execution by a
26 processor;
27 a window vector generating component, stored on the computer-readable
28 medium and executable by the processor, for mapping each database
29 sequence into a window vector; and
30 a tree-structure index generator, stored on the computer-readable medium and
31 executable by the processor, for generating a tree-structure index of the
32 window vectors.

33 22. The system of claim 21 wherein the k-tuple vector mapping component
34 further comprises:

35 a window partitioning component, stored on the computer-readable medium and
36 executable by the processor, for partitioning each database sequence into
37 a plurality of windows; and
38 a window vector mapping component, stored in the computer-readable medium
39 and executable by the processor for mapping each window into a window
40 vector.

41 23. A computer-implemented system for selecting database sequences from a
42 database comprising a plurality of sequences stored in electronic form, the system
43 comprising:

44 a computer processor for executing programmed instructions and for storing and
45 retrieving data;
46 an input device, coupled to the processor, for receiving a query sequence;
47 a program memory, coupled to the processor, for storing program instructions
48 for execution by the processor;
49 a tree-structure index generator, executable by the processor in accordance with
50 program instructions in the program memory, for generating a tree-
51 structure index that represents the database sequences;
52 a query sequence partitioning component, executable by the processor in
53 accordance with program instructions in the program memory, for
54 partitioning a query sequence into a plurality of windows; and

55 a retrieval component, executable by the processor in accordance with program
56 instructions in the program memory, for retrieving database sequences
57 from the database that are similar to the query sequence.

58 24. A computer readable memory, for use with a processor and a plurality of
59 database sequences, the memory storing a computer program executable by the processor
60 for configuring and controlling the processor to generate a tree-structured index that
61 represents the database sequences, by performing the steps of:

62 partitioning each database sequence into a plurality of overlapping windows;
63 for each window, generating a window vector representing the window; and
64 generating a tree-structured index of the window vectors.

65 25. A computer readable memory, for use in a computer system comprising a
66 processor and a database storing a plurality of database sequences in computer-readable
67 form, and for selecting database sequences from the database in response to a query
68 sequence, the computer readable memory storing:

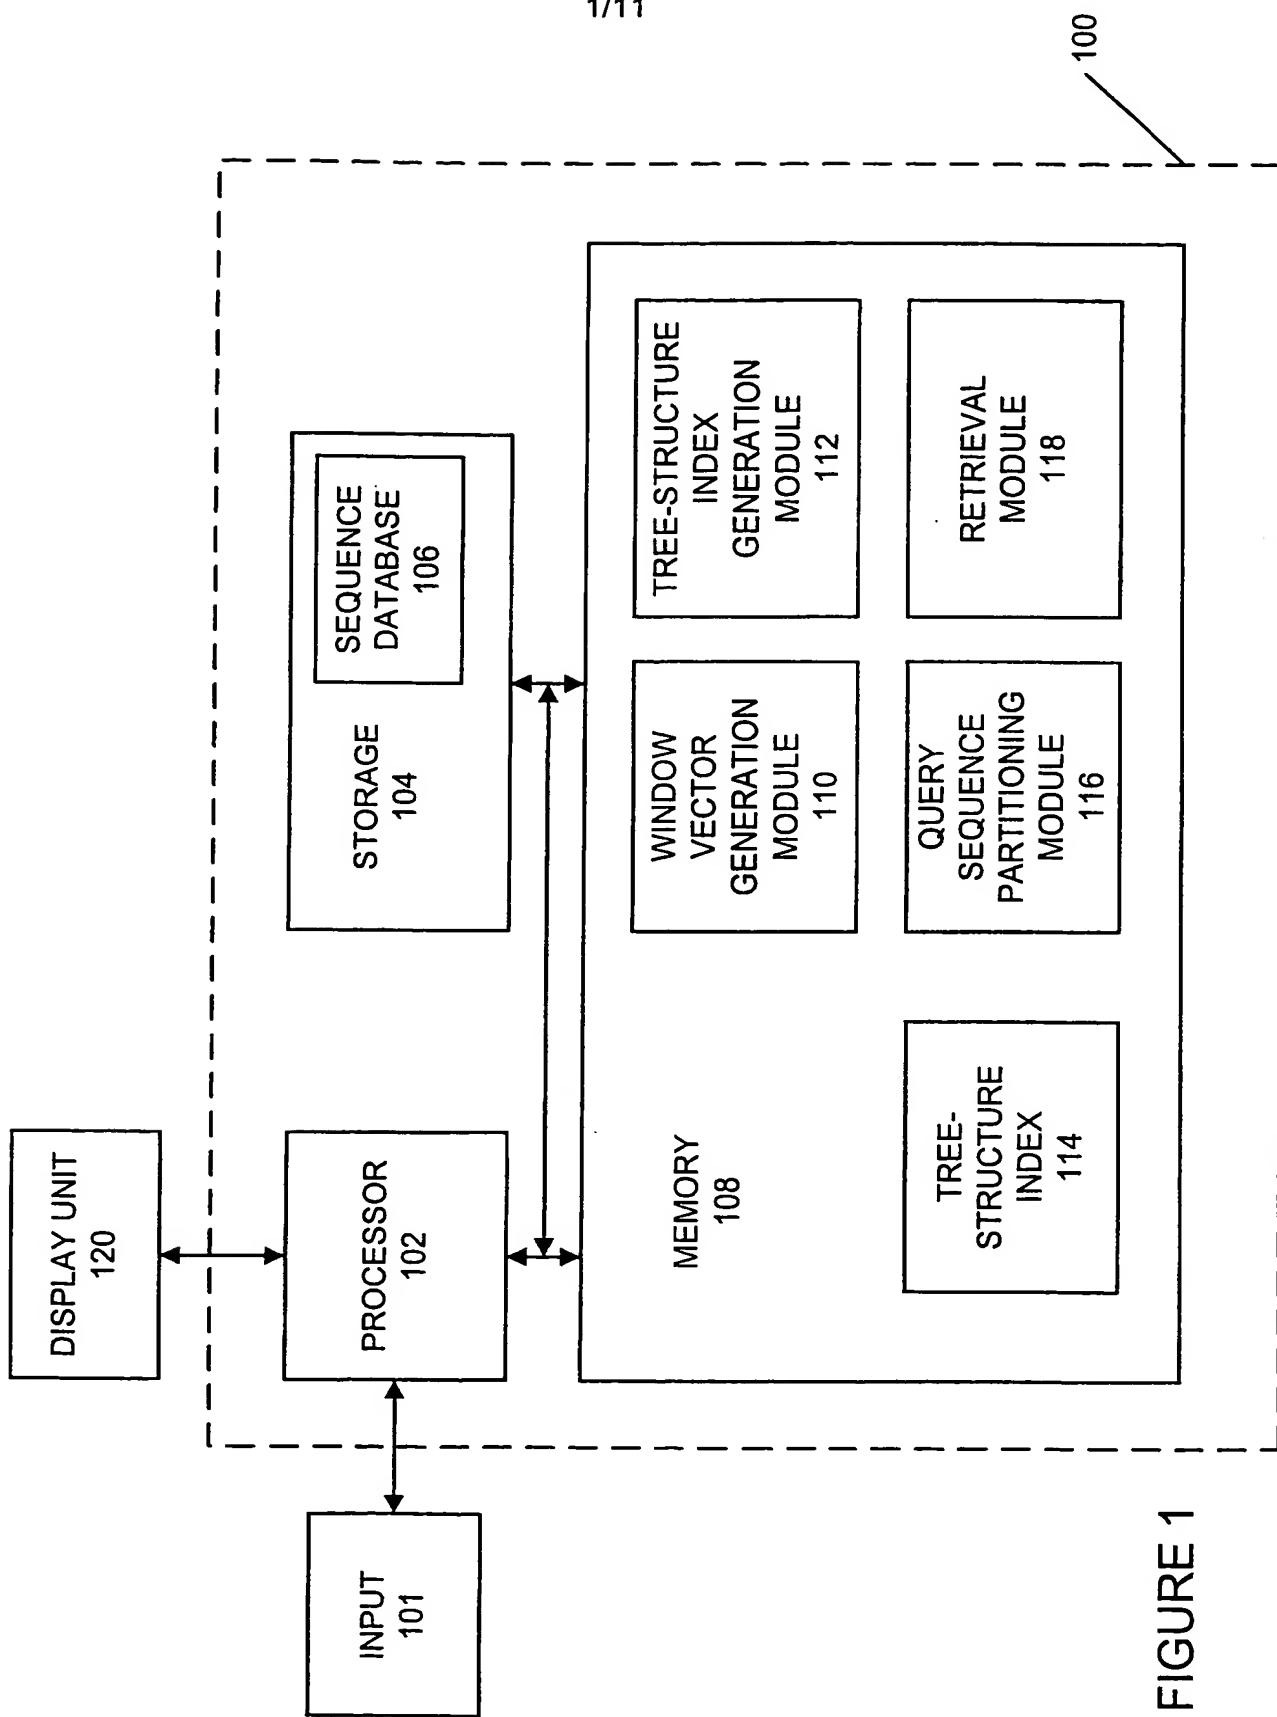
69 a tree-structured index generator, executable by the processor in accordance with
70 program instructions in the memory, for generating a tree-structured
71 index that represents the database sequences;

72 a query sequence partitioning component, executable by the processor in
73 accordance with program instructions in the memory, for partitioning a
74 query sequence into a plurality of windows; and

75 a retrieval component, executable by the processor in accordance with program
76 instructions in the memory, for retrieving database sequences from the
77 database that are similar to the query sequence.

78

1/11

**FIGURE 1**

2/11

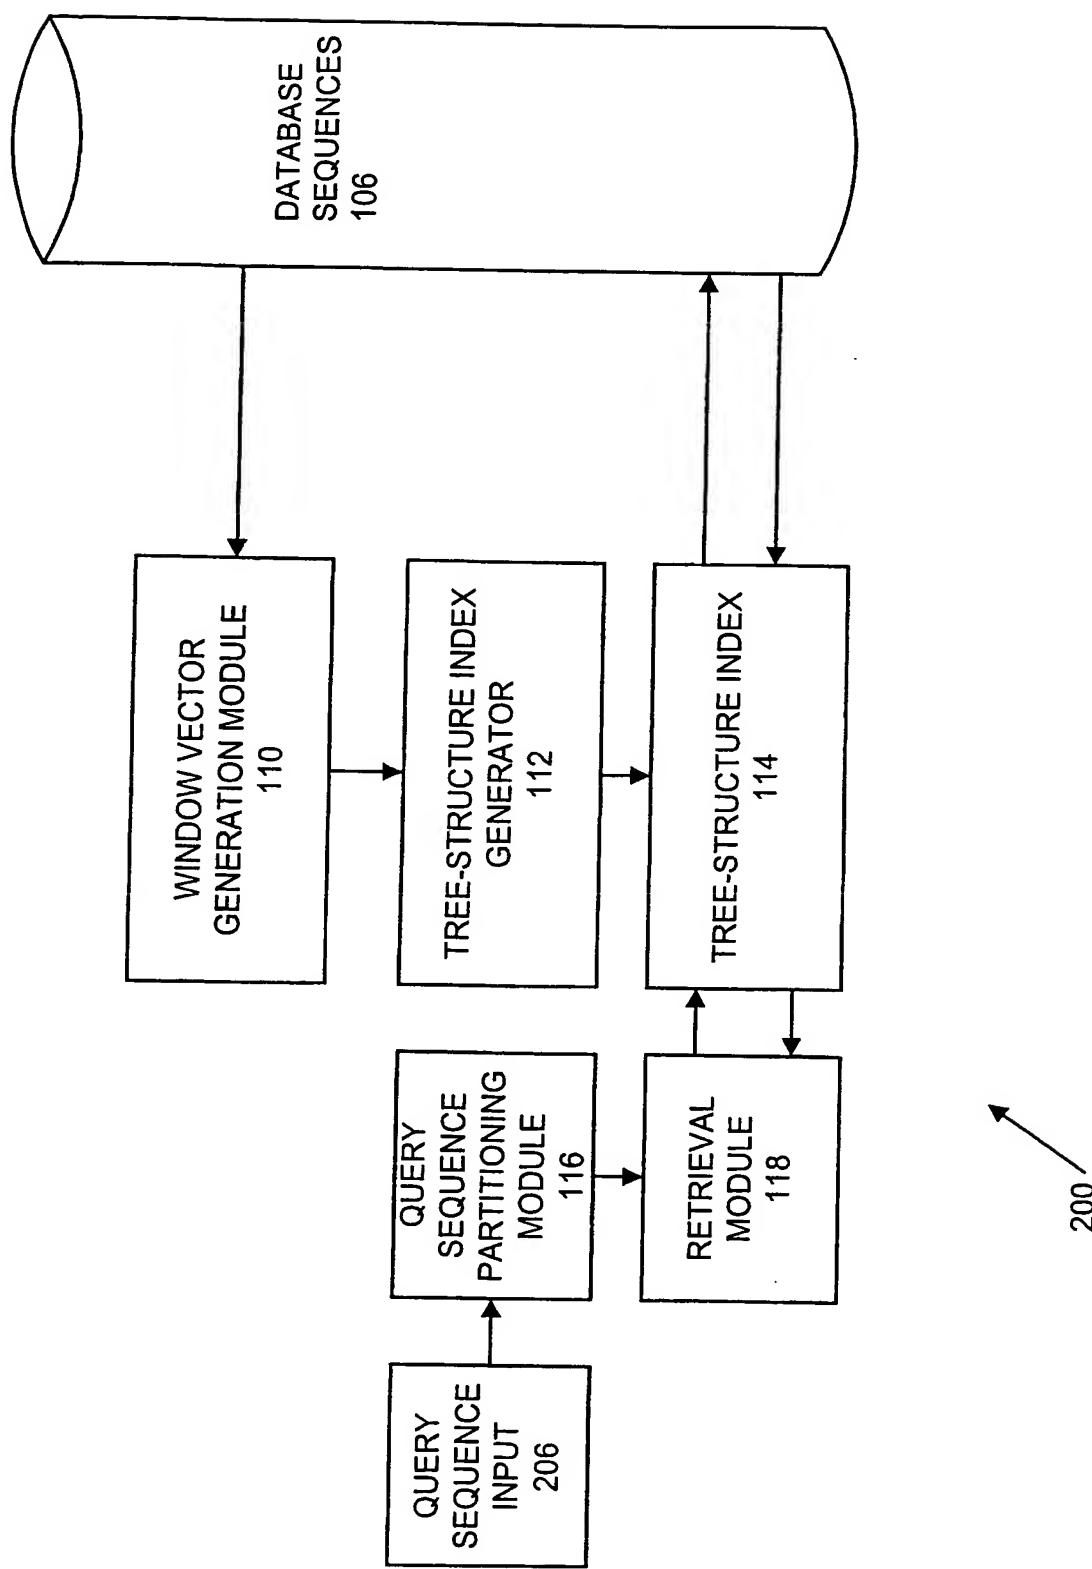


FIGURE 2

3/11

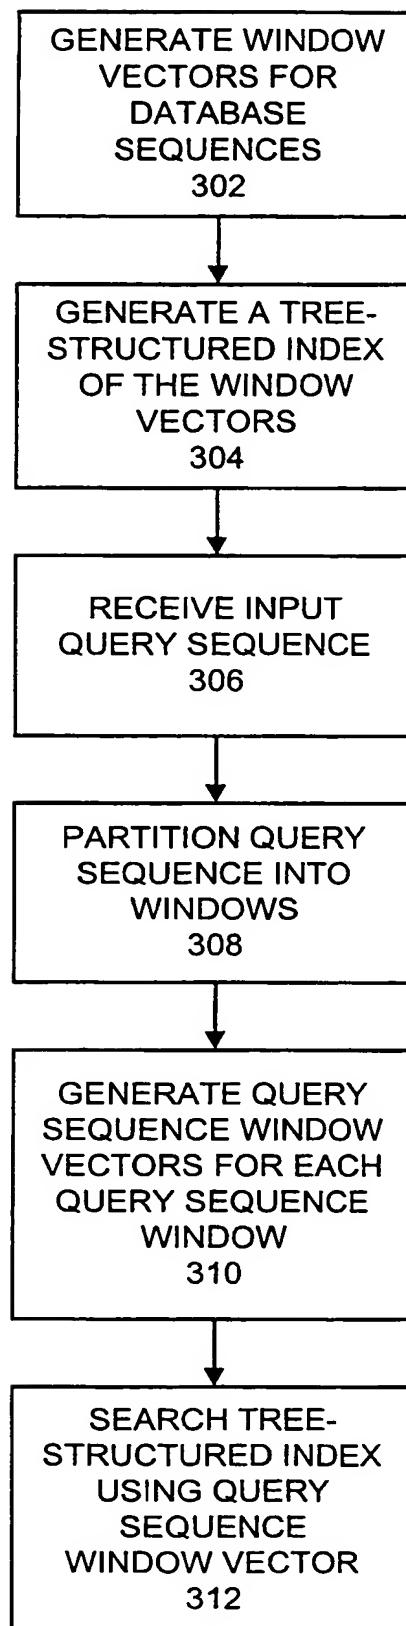


FIGURE 3

4/11

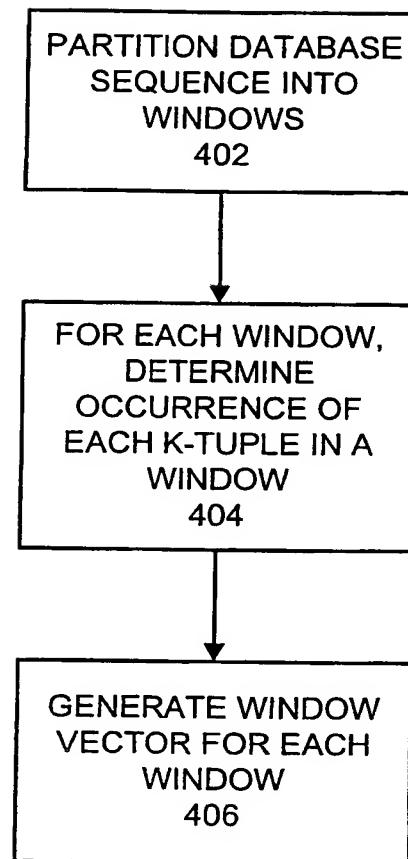
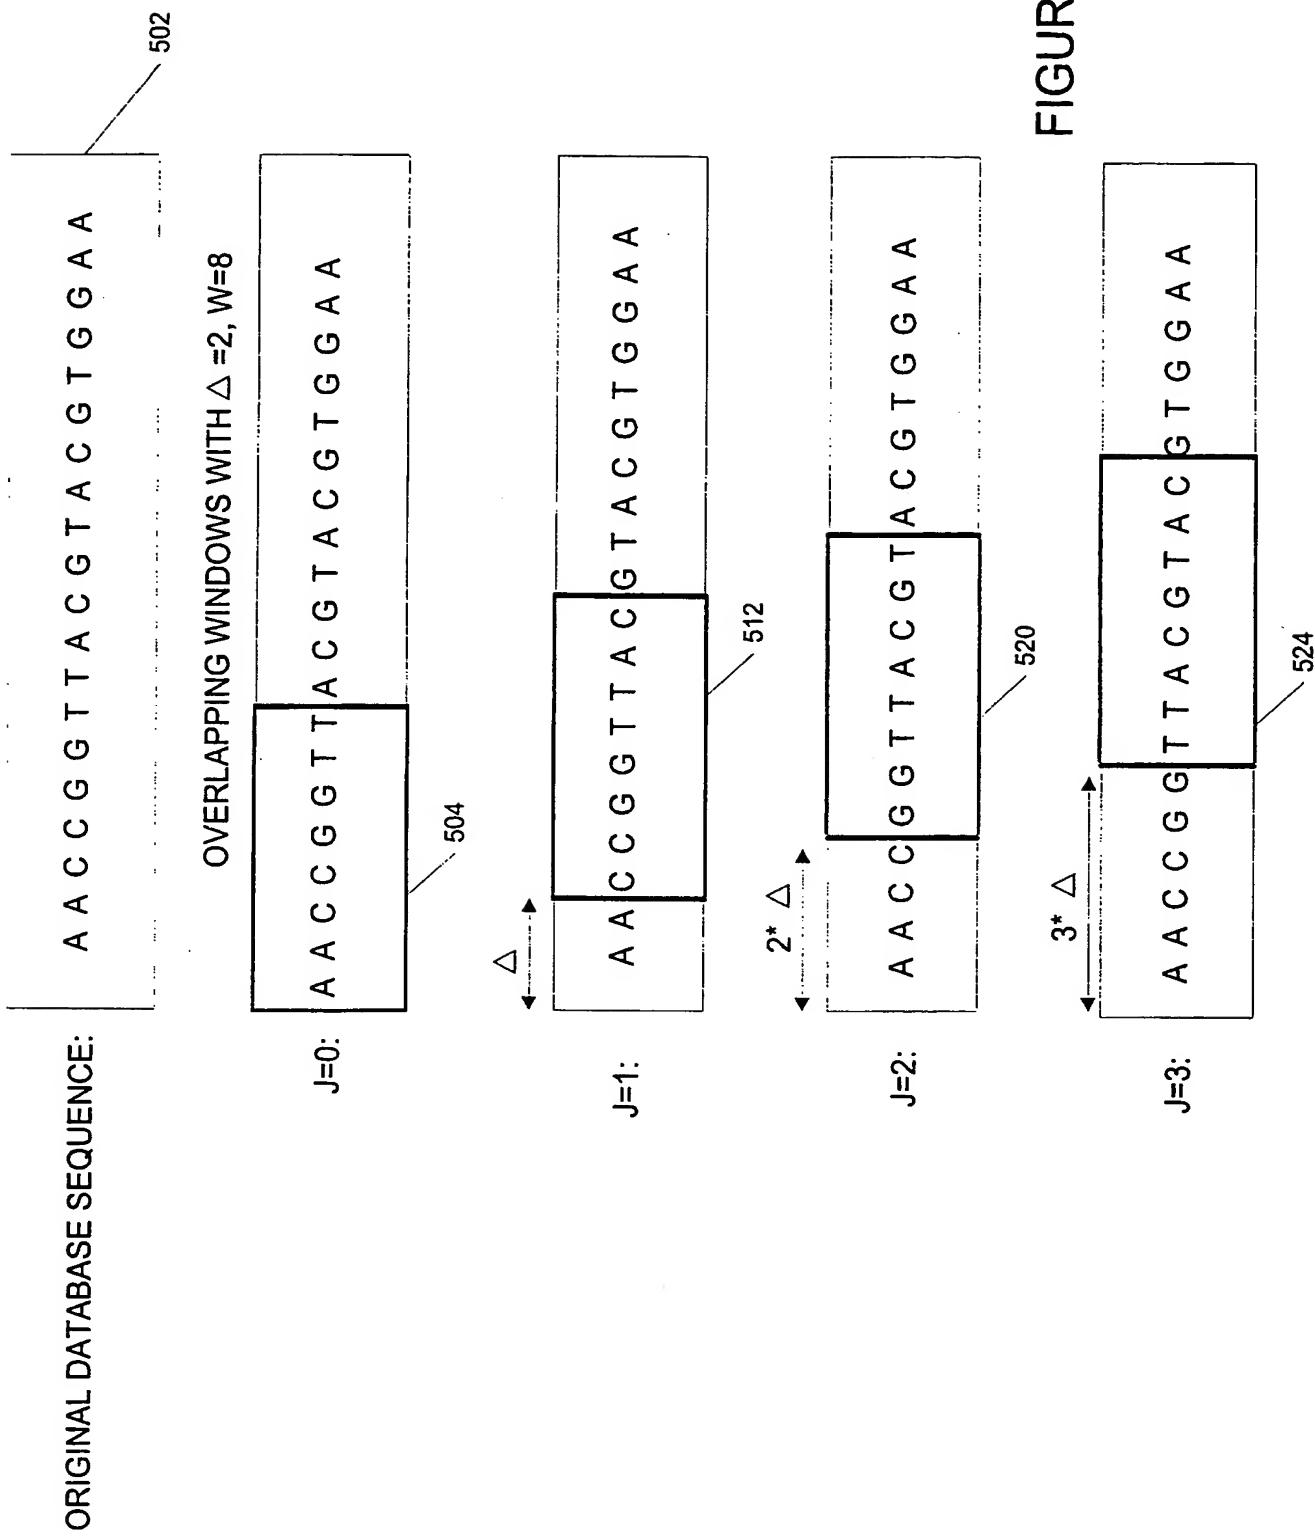


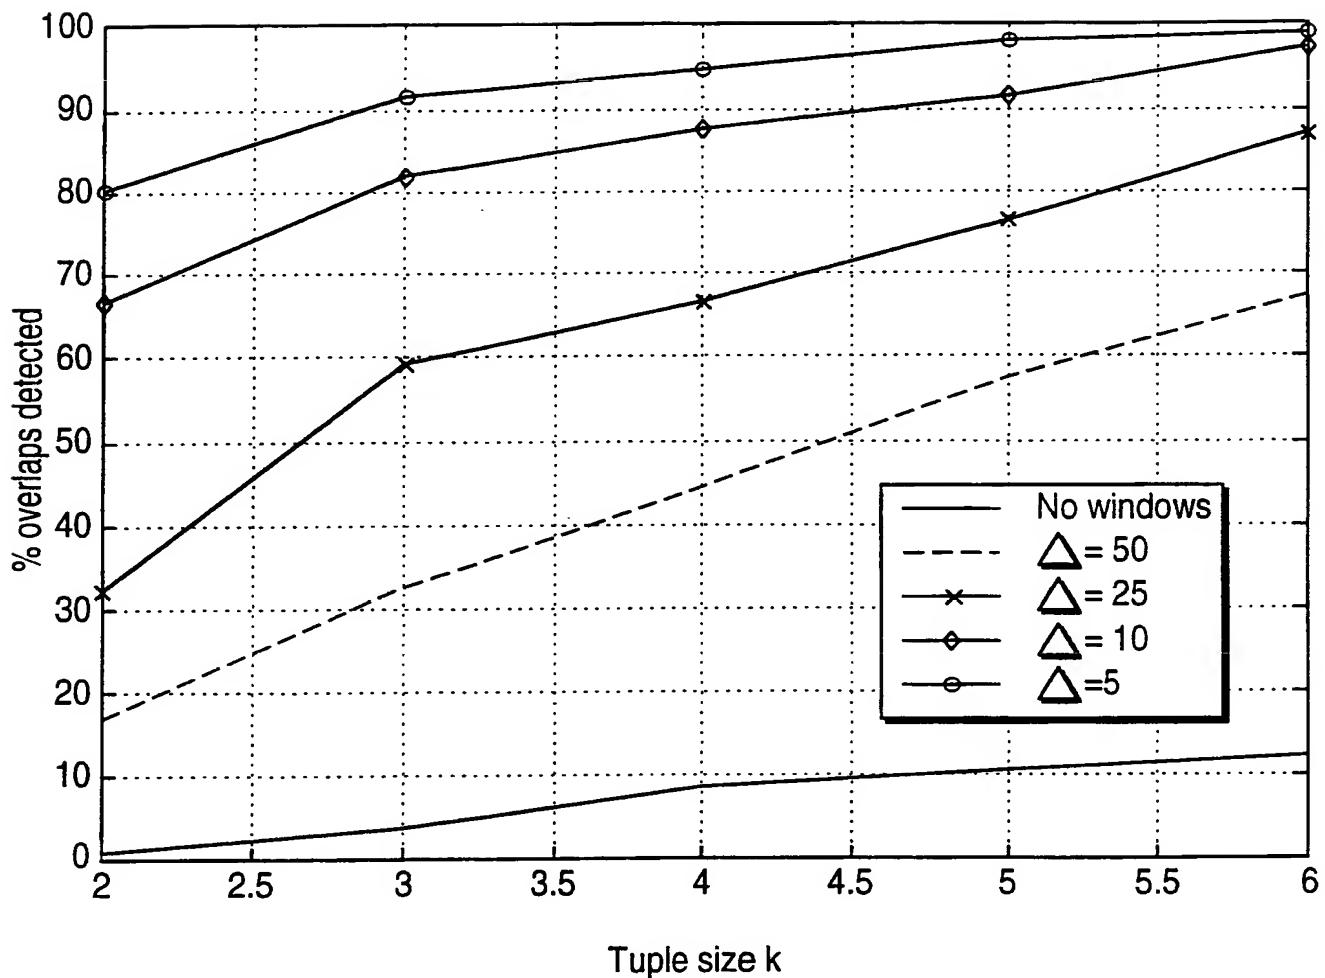
FIGURE 4

5/11



6/11

FIG. 6



7/11

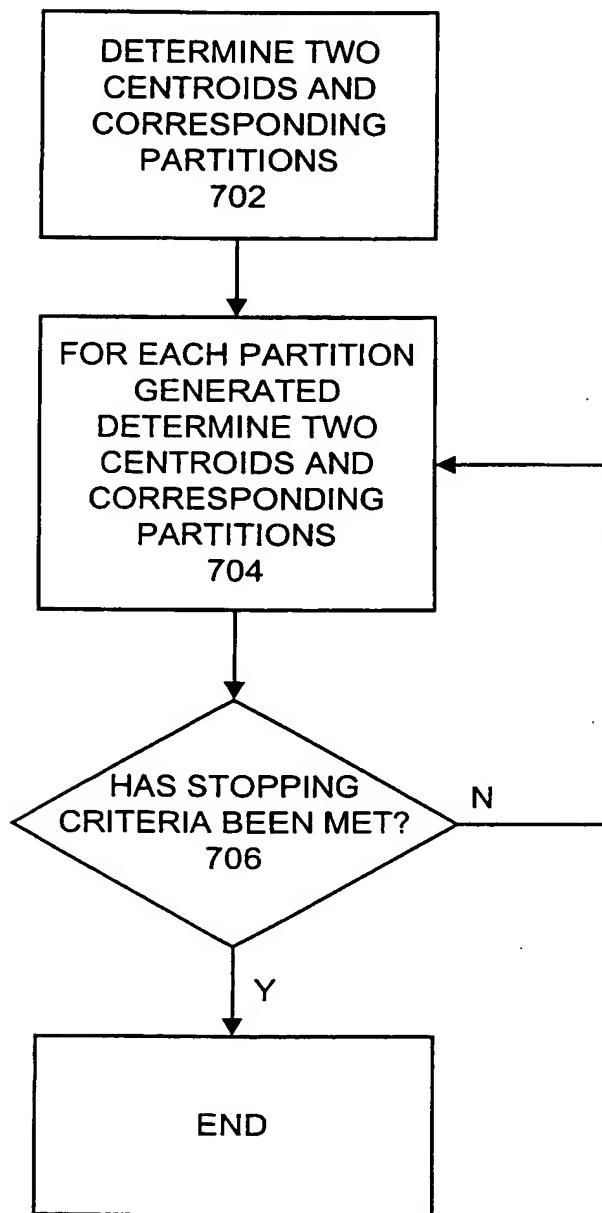


FIGURE 7

8/11

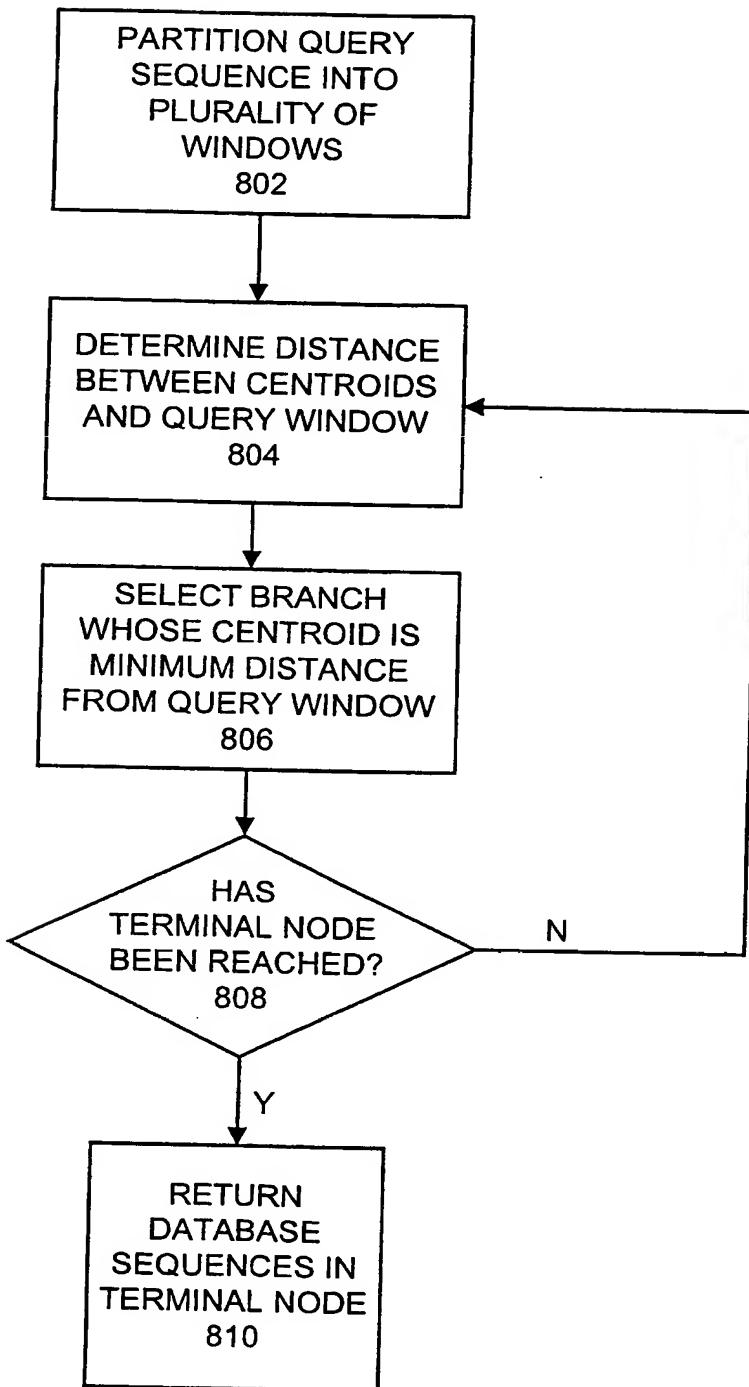


FIGURE 8

FIG. 9

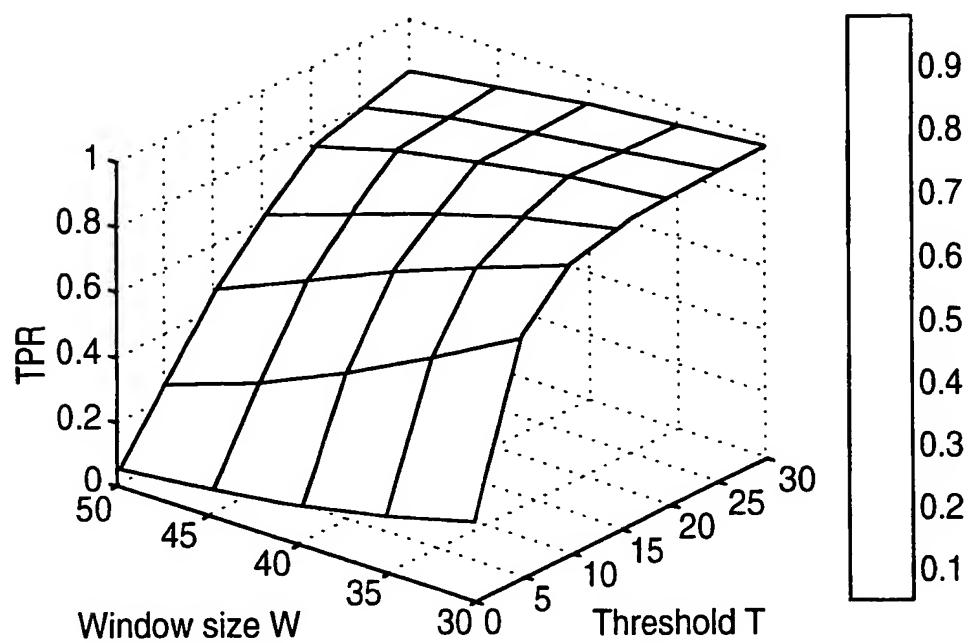


FIG. 10

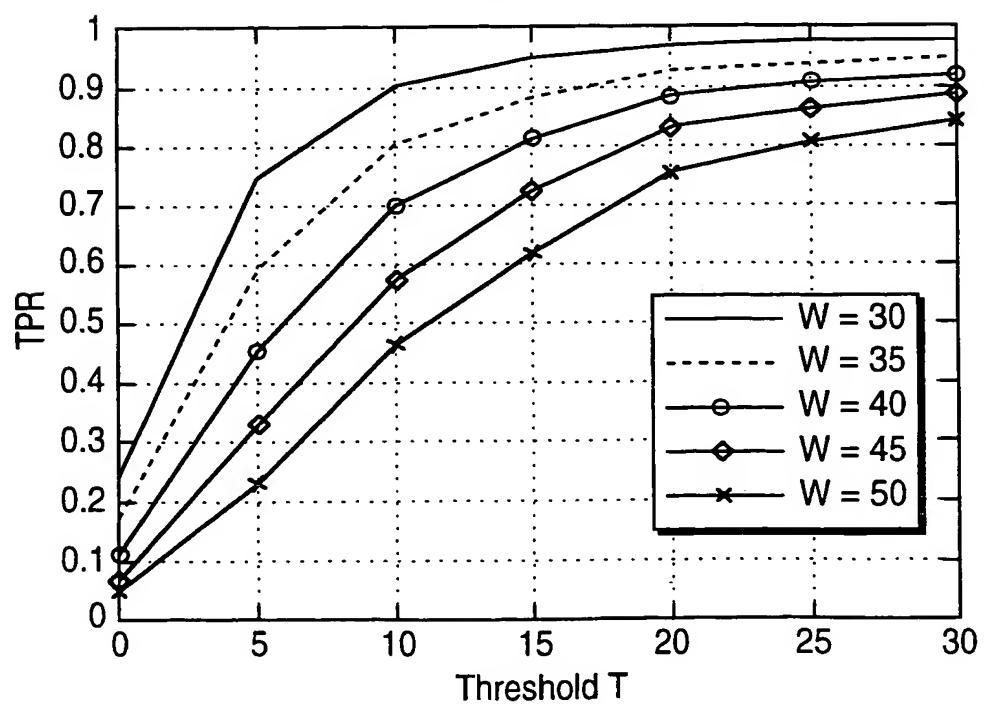


FIG. 11

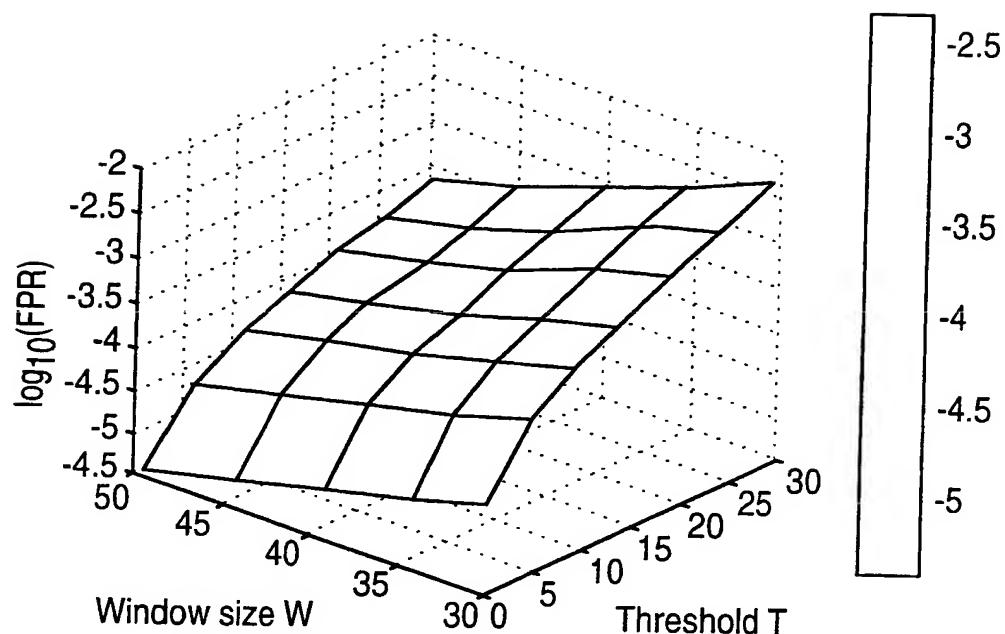
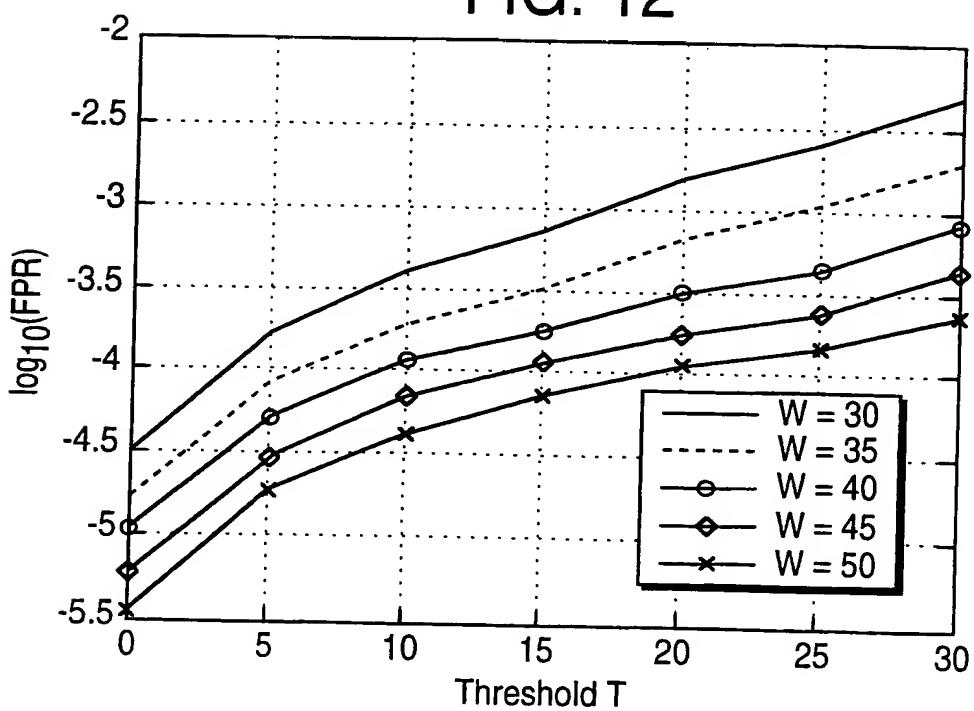


FIG. 12



11/11

FIG. 13

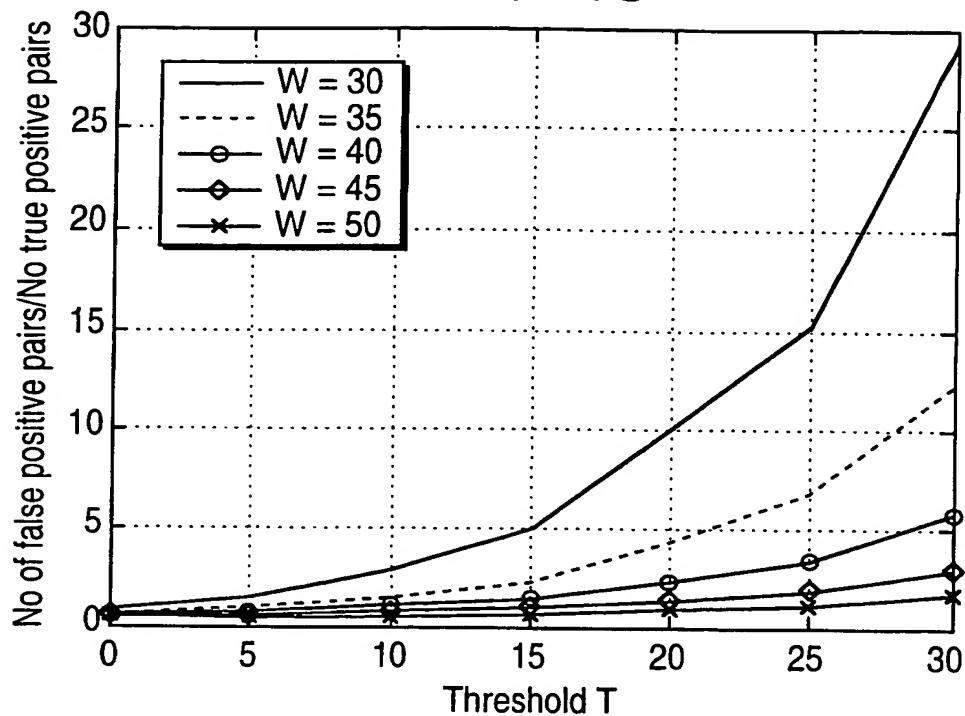
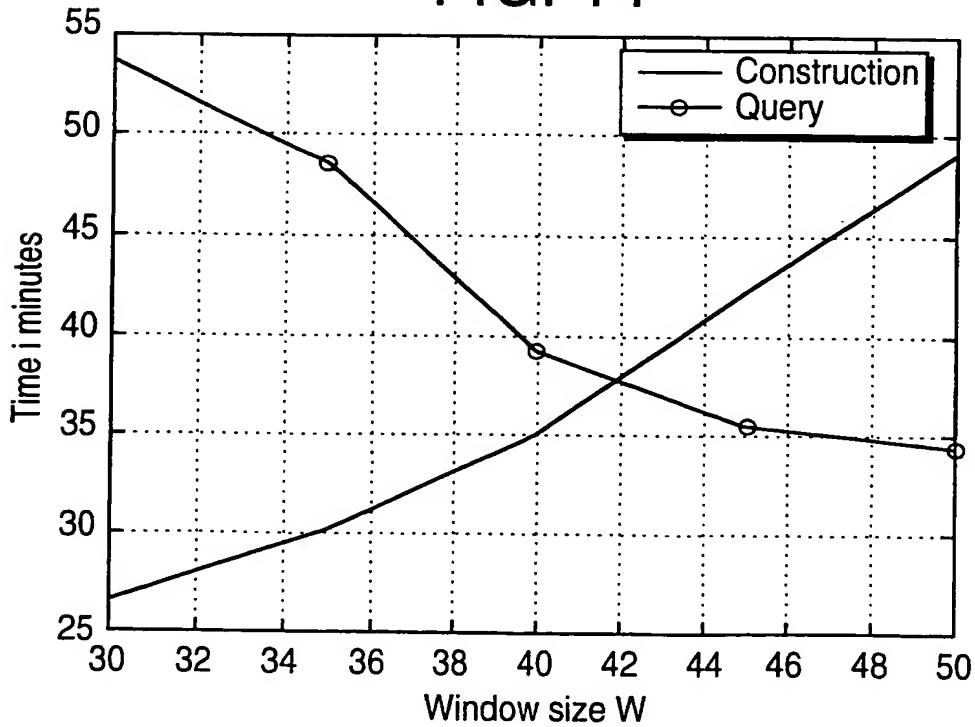


FIG. 14



INTERNATIONAL SEARCH REPORT

Internatinal Application No
PCT/US 00/34413

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F17/30

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F G06C C12Q

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, PAJ, WPI Data, INSPEC, IBM-TDB

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5 577 249 A (CALIFANO ANDREA) 19 November 1996 (1996-11-19) column 3, line 12 -column 4, line 63 ---	1, 10, 16, 21, 23-25
A	US 5 701 256 A (CHANG WILLIAM I-WEI ET AL) 23 December 1997 (1997-12-23) abstract column 4, line 42 -column 5, line 47 ---	1, 10, 16, 21, 23-25
A	EP 0 583 559 A (IBM) 23 February 1994 (1994-02-23) column 3, line 29 -column 4, line 37 --- -/-	1, 10, 16, 21, 23-25

Further documents are listed in the continuation of box C.

Patent family members are listed in annex.

* Special categories of cited documents :

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *&* document member of the same patent family

Date of the actual completion of the international search

23 April 2001

Date of mailing of the international search report

02/05/2001

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl.
Fax: (+31-70) 340-3016

Authorized officer

Deane, E

INTERNATIONAL SEARCH REPORT

Internat'l Application No
PCT/US 00/34413

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	PATENT ABSTRACTS OF JAPAN vol. 1996, no. 05, 31 May 1996 (1996-05-31) & JP 08 022475 A (MITSUBISHI ELECTRIC CORP), 23 January 1996 (1996-01-23) abstract -----	1, 10, 16, 21, 23-25

INTERNATIONAL SEARCH REPORT

Information on patent family members

Internat'l Application No

PCT/US 00/34413

Patent document cited in search report	Publication date	Patent family member(s)			Publication date
US 5577249	A 19-11-1996	EP	0583559 A		23-02-1994
		JP	2673091 B		05-11-1997
		JP	6098770 A		12-04-1994
US 5701256	A 23-12-1997	CA	2177764 A		01-12-1996
EP 0583559	A 23-02-1994	JP	2673091 B		05-11-1997
		JP	6098770 A		12-04-1994
		US	5577249 A		19-11-1996
JP 08022475	A 23-01-1996	NONE			